To Work Together: Examining the Time-Course of Multi-Effector Synergy in Motor Sequence Learning

William Kistler

University College London (UCL)

Department of Clinical and Movement Neurosciences

UCL Queen Square Institute of Neurology

A thesis submitted for the degree of

Doctor of Philosophy

December 2024

I, William Kistler, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

This thesis investigates the development of motor synergies during the acquisition of complex manual motor skills, specifically sequential typing tasks. Motor synergies - defined here as coordinated movements across multiple digits - form the foundation of skilled performance. Combining behavioural annotation, markerless pose estimation tools that estimate 2-D finger positions from video, and remote data collection methods, three studies were conducted to examine the emergence, refinement, and transfer of these coordinated movements during practice and rest in sequential motor learning.

Study 1 shows significant changes in motor synergies during early-stage practice. Over thirty-six trials, participants rapidly improved typing performance, with gains plateauing after early learning. Improvements coincided with more efficient, modular movement strategies, suggesting early establishment of control structures.

Study 2 shows that these patterns reorganise substantially during brief rest intervals rather than during active performance. This finding supports the idea that rest facilitates micro-offline consolidation of skill.

Study 3 explores transfer of coordination patterns from simple to complex tasks. Despite similarities in sensorimotor demands, coordination patterns from the simpler task did not transfer seamlessly to the complex one. Instead, participants formed new strategies to meet the demands of the five-element sequence.

Collectively, these studies advance our understanding of how motor synergies emerge and adapt during learning. They underscore the importance of both practice and rest, and the task-specific nature of motor control, offering insights for personalised training and rehabilitation.

Impact Statement

This thesis advances the scientific understanding of motor skill acquisition, with potential applications across clinical and technological domains. Motor sequence learning is foundational to skilled behaviour, yet traditional protocols often de-emphasise the importance of rest and the task-specific nature of motor coordination.

This research focuses on how motor synergies develop, refine, and transfer across tasks. It shows that significant consolidation of motor skill can occur over brief periods of rest, and that learned motor synergies are often specific to the practised task. These findings suggest that structuring rest with active practice could enhance motor learning - a principle relevant for motor and neuro-rehabilitation.

For example, future rehabilitation protocols could integrate structured rest with task-specific repetition to improve motor recovery in post-stroke populations, potentially accelerating reorganisation of coordinated motor activity. Similarly, in domains like surgery, athletics, or fine arts, strategically timed rest may help consolidate motor strategies more effectively than continuous practice.

Methodologically, this thesis introduces a scalable approach to motor behaviour assessment using markerless pose estimation, unsupervised clustering, and divergence metrics. Specifically, Jensen–Shannon Divergence (JSD), a measure of distance between distributions, was used to compare how the relative frequency of motor behaviours changed from one trial to another.

Future work could explore applying this framework in clinical trials or adaptive rehabilitation systems. By bridging neuroscience, machine learning, and rehabilitation science, this thesis contributes toward the development of more effective, individualised motor training strategies.

This thesis contributes to the broader goal of improving quality of life through better understanding of motor learning principles. By bridging gaps between neuroscience and rehabilitation, it lays a foundation for innovations that can enhance individual performance, reduce healthcare burdens, and improve societal well-being.

Acknowledgements

Thank you George**, Amanda, Justin, Jessica, Karen, Fran, Jay, and Becca for having inspired me. Thank you Meghan, Georgia, Dorothy, Chuck, Beth, Tim, Andy, Jaws, Jimmy, Nick, Oliver, Brendan, Jenny, Catharina and Mireia for inspiring me. Paul and Imran, thank you for your tireless support and for reminding me that help was always near. This effort was made possible with your passions and patience. My gratitude is outmeasured only by your generosity. Thank you.

Vielen Dank für Ihre Zeit, mein Freund.

Estoy eternamente agradecido por su guía.

Contents

List of Figures

List of Tables

Chapter 1

INTRODUCTION

"How did she do that?"

— GMM

1.1 Motivation for This Thesis

My daughter watched an Olympic diver twist smoothly through the natatorium air and enter the pool at La Défense Arena with only a short, foamy splash. "How did she do that," she asked, her eyes and heart fixed on the television. The dive impressed her enough that she needed to know. I couldn't fault her. I felt the same. "Honestly, love, I'm not sure". I avoided a complex reply and instead joined my daughter in her wonder. Of course we do know in part how athletic performance is formed at the behavioural, cognitive, and neurological level, but to her credit, many related questions in movement neuroscience have not been *fully* answered - how *did* she do that? How much practice must she have completed and how did she learn to blend each separate movement into one fluid execution? In this thesis, I will address each question by way of three experiments measuring kinematic and performance changes of participants learning novel keypress sequences.

* * *

This thesis is concerned with how humans move. Specifically, it addresses open questions regarding when and to what degree movement features emerge, change, or transfer over the course of learning new skills, during practice, and over periods of rest. In the next sections, I will introduce several topics relevant to the experimental chapters: motor skill learning, motor sequences, motor skill transfer, motor synergy, and the open questions surrounding motor synergy development. Each section will serve as a primer for the chapters to follow.

1. Motor Skill Learning and Practice

Motor skill development is fundamentally about learning through practice. From the earliest stages of learning - when an individual is consciously aware of each movement - to the later stages, where the movements become automatic, motor learning involves the constant refinement of internal models that guide action (Shadmehr and Krakauer, 2010; Krakauer and Mazzoni, 2019). In the case of the diver, each practice session would have involved refining the execution of individual movements as well as the integration of those movements into a coherent sequence.

Research in motor learning shows that practice is essential for skill acquisition because it allows the brain to update its internal models based on sensory feedback (Wolpert and Flanagan, 2001). Over time, these updates help the learner improve the timing, accuracy, and efficiency of their movements. This principle applies to elite athletes as well as everyday motor tasks like learning to walk or type. As movements are practised and refined, they shift from being explicitly controlled to becoming more implicit, relying on the brain's ability to execute actions without conscious oversight (Adams, 1971; Shadmehr and Krakauer, 2010).

2. Motor Sequence Learning: Combining Movements

One of the key processes behind motor learning is motor sequence learning - the ability to combine individual movements into a continuous sequences. In the diver's case, this means knowing not just how to execute a dive but also how to transition smoothly from one movement to the next. Motor sequence learning is essential for any skilled task that involves multiple steps, from playing a musical piece to

performing a complex athletic routine (Abrahamse et al., 2013; Karni et al., 1995).

Initially, learners master individual components of a sequence, consciously planning each action. But with practise, these individual actions become "chunked" into larger units, making it easier to recall and execute entire sequences (Verwey, 2001). For the diver, this means that the entire dive - composed of various twists and turns - can be executed as a single fluid motion rather than a series of disjointed movements. Research suggests that as motor sequences are learned, the brain improves prediction of outcomes of movements, allowing for smoother transitions between actions (Hikosaka et al., 2002).

3. Motor Synergies: Coordinating Movements

Underlying the fluidity of the diver's movements is another important concept: motor synergy. This refers to the way the nervous system coordinates muscles, joints, or movements to produce smooth, efficient movements (Latash, 2008a). With over 600 muscles in the human body, coordinating them all to produce a simple movement, let alone a complex dive, is challenging. Motor synergies simplify this coordination by grouping muscles movements together into functional units, allowing efficient and effective performance (Bizzi et al., 1991; Turvey, 2007).

For the diver, motor synergies ensured muscle groups works together to produce the mastered movements with minimal effort. This coordination is not something that happens overnight; it is the result of extensive practice, where the body learns to exploit synergies to minimize effort and maximize performance. As these synergies become more refined, the diver is able to execute increasingly complex dives with greater precision and control (Bernstein, 1967; Latash and Huang, 2015).

4. Motor Skill Transfer: Solving Redundancy

Motor skill transfer refers to the process where learning one motor skill influences the acquisition of another. This phenomenon can occur when the components of a newly learned skill share similarities with a previously mastered skill (Schmidt, 1975; Wolpert et al., 2011). When the Olympic diver, for example, executes a new and complex dive, much of their ability to quickly master it comes from previ-

ously learned dives that share similar motor components. In essence, the nervous system identifies and implicitly transfers muscle activation patterns and joint coordinations, thereby minimising the need to re-learn redundant motor behaviours (Berniker and Kording, 2008; Diedrichsen et al., 2005).

Solving redundancy in motor control is central to transfer. Redundancy refers to the fact that there are often multiple ways to achieve the same motor goal, with different combinations of muscle activations and joint movements producing the same result (Latash, 2010). To solve for computational redundancy, the body may reuse previously cost-efficient and successful patterns when a new skill incorporates elements of those previously learned. In the case of our diver, this means that muscle synergies and coordination strategies from earlier dives aid in learning new dives. The reuse of motor components underpins the concept of motor skill transfer and is a critical factor in motor learning research though large gaps remain in its understanding (Newell, 1991; Turvey, 1990).

5. Open Questions: The Development and Transfer of Synergy

In the field of motor learning, there are many open questions surrounding motor synergy development and their timescale - how long does it take for motor synergies to develop fully, and once developed, can synergies transfer to new tasks?

While research has shown that motor synergy is crucial in motor sequence learning, the exact mechanisms and timescale by which they develop and transfer remain unclear. Understanding these processes could have substantive implications for athletics and rehabilitation, where the ability to transfer motor skills from one context to another is critical. This thesis will address how and when motor synergy emerges, evolves, and transfers amid motor sequence learning.

1.2 Motor Skill Learning

Motor learning can be defined as any experience-dependent improvement in motor performance (Wolpert et al., 2011; Krakauer and Mazzoni, 2019). As highlighted earlier, motor learning is crucial not only for expertise in activities like athletics but also for day-to-day tasks. From a child learning to walk or tie their shoes to an

adult learning to drive or an elderly person learning to use new technology, motor learning underpins many essential life skills. Moreover, it is especially relevant for patients undergoing rehabilitation following motor disabilities such as stroke, where re-learning skills becomes critical for functional recovery (Krakauer et al., 2000).

Daily motor tasks are incredibly diverse - whether diving or using the pedals while driving - and likely engage various learning mechanisms. One point of agreement in the field is that motor learning typically results in improvements in the speed and/or accuracy of movements, ultimately enhancing the speed-accuracy tradeoff (Fitts, 1954). For example, a tennis player may improve by increasing both movement initiation speed (e.g., allowing faster reactions and more powerful strokes) and movement accuracy (e.g., choosing the correct stroke and executing it reliably). Other aspects of motor learning include reduced susceptibility to external perturbations (e.g., maintaining movement accuracy under force perturbations) and the ability to perform tasks with reduced cognitive load, which can be measured using dual-task paradigms (Johansson et al., 2023). These phenomena indicate that motor learning occurs at multiple levels and is modulated to the task's demands.

In laboratory settings, motor learning is typically studied through a variety of tasks, which provide insights into the mechanisms underlying this process. Two dominant paradigms have emerged: motor sequence learning and motor adaptation (Krakauer and Mazzoni, 2019). In motor sequence learning tasks, participants are asked to perform a series of key presses on a keyboard or button box as quickly and accurately as possible, emphasizing learning of successive movements. In contrast, motor adaptation tasks require individuals to adjust a well-learned movement (such as a reaching task) in response to external perturbations, focusing on correcting and adapting movements. See a brief table summarising the differences between sequence learning and adaptation in Table 1.1. These paradigms have advanced our understanding of the neural substrates involved in motor learning. However, there is ongoing debate about whether they serve as optimal models for real-world skill acquisition (Doyon et al., 2003; Doyon and Benali, 2005; Morehead and Orban de Xivry, 2021). A key difference between these tasks and real-world skills lies in the

emphasis on movement execution. While laboratory tasks often prioritise movement selection (e.g., selecting the correct sequence of key presses or correcting for a perturbation), real-world tasks like tennis require both accurate selection and precise execution of movements (Haar and Faisal, 2020).

Table 1.1: Comparison of Motor Sequence Learning and Adaptation Tasks

Feature	**Sequence Learning**	**Adaptation**
Primary Goal	Learn a novel sequence	Compensate for perturbation
Example Task	Serial Reaction Time, typing	Force-field reaching
Neural Involvement	Striatum, SMA, M1	Cerebellum, M1
Learning Outcome	Faster, more accurate execution	Realignment of motor commands
Consolidation Sensitivity	Sleep-dependent	Rapid, error-driven

In this thesis, I focus on the former paradigm, motor sequence learning, and I model learning via a sequential typing task, where participants executed a series of key presses as quickly and accurately as possible. This task captures the temporal and spatial ordering of motor elements, requiring the participants to recall and perform sequences with precision and speed. By emphasizing motor execution, the typing task provides an excellent model for understanding how motor sequences are structured, learned, and refined over time. The choice of this task reflects its relevance to daily life activities like typing or playing a musical instrument, both of which rely on fine motor control and the development of complex motor synergies.

Additionally, this task enables the exploration of the motoric contributions to sequence learning - specifically, the development of motor synergies and their timescale. Investigating the formation of motor synergy through a sequential task allows us to gain insight into how the brain organizes and optimizes motor commands to produce consistent, coordinated movements. The study of motor synergies is crucial for understanding the underlying processes that make motor learning efficient and flexible, as they allow the motor system to manage the complexity of movement by combining multiple motor elements into cohesive units.

While the focus of this thesis is on motor sequence learning, it is important to situate such learning within the broader context of how practice influences be-

haviour. Recent theoretical work suggests that practise leads to at least three distinct but interrelated behavioural changes: increased skill (improved speed and accuracy), reduced cognitive load, and the development of habitual responses (Haith and Krakauer, 2018). These effects, sometimes collectively referred to as automaticity, do not necessarily evolve in lockstep and may reflect caching of different intermediate computations. Motor sequence learning, as studied here, likely engages a proportion of all three dimensions - requiring initial effortful selection and execution of actions, which with practice become faster, less attentionally demanding, and potentially habitual. Exactly which components of a manual keypress task are governed more by skill or habit as they are defined by Krakauer and Haith - e.g. the keypress itself or the transition between any two keypresses - is topic up for debate. A richer understanding of motor skill learning thus benefits from considering these dissociable but co-occurring effects of practice.

This thesis, therefore, focuses on the cognitive and motoric aspects of sequence learning, investigating how motor synergy develops over time and whether these synergies can be transferred across skills. This focus enables an expanded understanding of motor learning processes, as they apply both in controlled laboratory settings and in real-world scenarios.

1.3 Motor Sequence Learning

Considerable research in motor learning over the past century has focused on how actions are organized into specific temporal sequences to achieve tasks. This organization may involve discrete movements, such as those needed to prepare tea, or continuous and overlapping actions within a single movement, such as muscle activations for a tennis serve (Abrahamse et al., 2013; Ahmadi-Pajouh et al., 2012; Berniker and Kording, 2011). Understanding how sequences are learned, reproduced, and represented in the brain is crucial for insights into motor control.

In daily activities, task completion also necessitates performing a sequence of actions in a prescribed order. For example, making tea involves actions such as filling a kettle, boiling water, placing a teabag in a cup, and pouring water over

it. While certain steps can be reordered without significantly affecting the outcome (e.g., placing the teabag before or after adding water), others must follow a strict sequence (e.g., boiling water before pouring it) (Benke, 1993). Thus, the goal of the sequence is to complete the task, with individual actions well understood.

Language is another domain where sequence learning is fundamental. Rules of phonology, morphology, and syntax dictate the order in which sounds and words are arranged to convey meaning (Lashley, 1951). These rules, in turn, influence motor control of speech production. In both motor behaviour and language, individual elements often become grouped into chunks - hierarchical groupings that facilitate the representation and recall of subsequences (Jeannerod, 1988; Albouy et al., 2008). For example, a sequence of sounds can form a word that is executed through a specific pattern of muscle movements. Patients with apraxia of speech struggle with producing these sequences, resulting in inaccurate sound production (Benke, 1993). In laboratory settings, action sequences are often studied using simplified tasks such as button-pressing sequences or saccadic eye movements, with an emphasis on learning to select and execute actions in the correct order (Verwey, 2001). Over time, these actions are generated more automatically and accurately, reinforcing the notion of motor chunking (Abrahamse et al., 2013).

Motor sequence learning can be broadly categorized into two types: learning discrete action sequences and learning continuous movements. Discrete sequence learning focuses on the correct temporal order of distinct actions, while continuous actions involve overlapping or uninterrupted movements, such as a tennis serve, where task goals are achieved through the coordination of multiple muscle activations (Abrahamse et al., 2010; Albouy et al., 2013). Although these two forms of learning are distinct, they are often studied together. For example, a patient with ideational apraxia may be able to perform a simple reaching movement but struggle with completing a series of discrete tasks (Benke, 1993). This distinction highlights the potential involvement of different motor planning stages during improvements in sequential tasks (Abeele and Bock, 2001).

In discrete sequence tasks, the goal is to rapidly and accurately select and ex-

ecute each action in a prescribed order. For instance, in the Discrete Sequence Production (DSP) task, participants practice keypress sequences, improving speed and accuracy as the movements become more automatic (Abrahamse et al., 2013). Over time, as finger movements overlap temporally, sequence execution becomes faster and smoother, supported by the formation of new neural representations that underlie this performance improvement (Abdelghani et al., 2008). Similarly, the Serial Reaction Time Task (SRTT) involves executing movements in response to stimuli presented in a fixed sequence. This task allows researchers to assess sequence learning by comparing response times for learned sequences versus random sequences (Albouy et al., 2008).

The second conceptualization of sequence learning, involving continuous actions, focuses on refining the execution of uninterrupted movements, such as a tennis serve or a simple reach. These tasks involve the sequential activation of muscles in a precise order, coordinated through complex motor commands. For example, reaching movements are associated with the triphasic burst pattern in fast-reaching tasks or the coordinated activation of extraocular muscles during saccades (Ramat et al., 2007; Robinson, 2022). Unlike discrete sequences, the motor system generates these patterns unconsciously, reflecting the smooth and continuous nature of the movement (Ahmadi-Pajouh et al., 2012).

In this context, while discrete sequence learning provides insights into how the brain encodes the order of individual actions, and continuous movement studies highlight the complex coordination of muscle activations required for fluid execution, tasks that embody elements of both paradigms offer a unique opportunity to explore the interplay between discrete and continuous motor learning mechanisms. In this thesis, I use a typing task that integrates both concepts. Although the task involves discrete actions - individual keypresses - it demands that the whole sequence be executed repetitively without pauses, mirroring the characteristics of continuous movements. This design allows us to investigate how the motor system transitions from planning discrete movements to achieving coordinated, continuous execution. A key mechanism that facilitates this transition is the formation of motor chunks,

where individual actions are grouped into cohesive units.

1.3.1 The Formation and Representation of Chunking

In motor sequence learning, a key characteristic is that once a sequence has been mastered, its individual elements are no longer executed independently but are grouped into collective units, or chunks. This grouping forms a more efficient neural representation of the entire sequence, allowing for rapid recall and execution. Evidence supporting this idea comes from the observation that the reaction time (RT) for the first movement in a sequence increases as the number of subcomponents increases, suggesting that the entire ordered set of actions is recalled before initiating the first element (Verwey, 2001; Abrahamse et al., 2013). Instead of planning each movement separately, the motor system treats the sequence as a unified pattern, storing subsequent action order in a buffer for efficient execution (Verwey, 2001; Hikosaka et al., 2002).

This principle of chunking is not unique to the motor domain: expert memory performance in complex cognitive tasks, such as chess, similarly relies on the rapid retrieval of pre-formed knowledge structures. For example, Gobet and Simon demonstrated that chess masters could recall multiple board positions far beyond typical short-term memory limits by drawing on domain-specific templates - essentially chunked configurations of pieces - stored in long-term memory (Gobet and Simon, 1996). Such cross-domain parallels suggest that chunking reflects a general mechanism by which the brain simplifies complex information into coherent, retrievable units.

However, many motor sequences can be quite long and complex, and evidence suggests there is a limit to the number of elements that can be grouped into a single unit. Typically, sequences are divided into smaller chunks, each containing 3-7 elements (Miller, 1956; Verwey and Eikelboom, 2003). These “chunks” are often indicated behaviourally by an increase in response time for the first element of a new chunk, suggesting that the motor system is recalling and pre-planning the entire chunk at once (Verwey, 1996; Verwey and Eikelboom, 2003). For example, in keypress sequences, participants typically show slower response times at the start

of each chunk, followed by faster responses for the remaining elements. This pattern reflects the cognitive grouping of actions into chunks, which facilitates more efficient sequential execution (Verwey, 1996; Abrahamse et al., 2013)..

Recent evidence, however, has suggested that these changes in response time may not always correspond to chunk boundaries, but might also reflect biomechanical constraints or idiosyncratic strategies used by participants (Ramkumar et al., 2016). Despite these challenges, the idea of chunking remains central to how sequences are organized in the brain. Chunking allows for sequences that exceed working memory capacity to be recalled more efficiently, by reducing the cognitive load to remembering a few chunks rather than many individual elements (Verwey, 2001; Verwey and Eikelboom, 2003). This strategy is not unique to motor sequences but is also seen in language and problem-solving, where individuals rely on chunking to manage large amounts of information (Chase and Simon, 1973).

The specific length and structure of chunks may be influenced by working memory capacity and the degree to which a sequence can be compressed or simplified into a more manageable form (Miller, 1956; Verwey, 2001). For instance, in skilled performance, such as chess or music, experts often utilize chunking strategies to recall complex configurations or sequences that would otherwise overwhelm working memory (Chase and Simon, 1973). This ability to form and manipulate chunks is thought to be a fundamental cognitive strategy that supports both motor learning and broader cognitive functions.

The chunking process typically begins with the associative learning of simple transitions between small fragments of a sequence, such as pairs of responses (Johnson, 1970; Miller, 1956; Laird et al., 1984; Graybiel, 1998; Servan-Schreiber and Anderson, 1990; Verwey et al., 2015). Initially, these transitions are stimulus-dependent, but with practice, they evolve into longer, more stable fragments that can be executed with increasing autonomy (Verwey et al., 2015). Over time, these chunks become reinforced through repetition, resembling statistical learning mechanisms (Perruchet and Pacton, 2006; Hunt and Aslin, 2001).

Chunking provides a distinct performance advantage. Sequences constructed

by recombining familiar chunks are executed more efficiently than completely novel sequences, though not necessarily as fluently as well-practised ones (Wymbs et al., 2012). This suggests that not only are individual elements learned within chunks, but the transitions between them are also integrated as part of the sequence representation (Diedrichsen and Kornysheva, 2015). Moreover, chunking emerges dynamically during skill acquisition: practice strengthens these structured representations, but inefficient chunking patterns established early in learning tend to persist despite extensive training (Verwey et al., 2015). This suggests that chunking is not solely a mechanism for optimising motor execution but rather a process of structuring and organizing movement elements into manageable units.

The formation of chunks is task-dependent: advantages arise when an entire sequence is practised as a unit rather than through isolated repetitions of common sub-elements (Kornysheva and Diedrichsen, 2014; Kornysheva et al., 2019). However, the extent to which chunking contributes to long-term memory formation and consolidation remains uncertain. While hierarchical and modular structures of motor sequences appear to stabilise with practice, it is still unclear whether these structures strengthen through consolidation or whether they are simply a function of initial learning organization (Diedrichsen and Kornysheva, 2015).

There has been significant debate regarding whether higher-order representations, like chunks, are motoric or exist at a cognitive level. Lashley (1951) proposed that higher-order representations organize sequence elements abstractly rather than focusing on the motor details of individual actions. This view is supported by findings that sequences learned symbolically, such as through colour-coded cues, can transfer across different effectors, including saccadic eye movements (Verwey et al., 2015). Similarly, sequence learning through observation can yield performance benefits comparable to active practice, suggesting that chunking can emerge independently of direct motor execution (Bird and Heyes, 2005; Hardwick et al., 2018). The ability to transfer sequence rules, where the specific elements change but their relationships remain intact, also points toward a hierarchical, cognitive organization that is action-independent but order-dependent (Schmidt and Lee, 2005).

Recent neuroimaging studies bolster the notion that chunking is hierarchical. Activity in the primary motor cortex (M1) during motor sequence learning does not necessarily reflect a reorganization of individual movement representations but rather a compressed encoding of sequential dependencies (Wymbs et al., 2012). Even after extensive training, motor chunking appears to be represented more robustly in premotor and parietal areas than in M1, supporting the idea that chunking is a strategy for organizing sequence order rather than transforming motor execution itself (Tanji and Shima, 1994; Kennerley et al., 2004; Kornysheva et al., 2019). Cortico-striatal interactions, particularly within the basal ganglia, have also been implicated in sequence chunking, with evidence suggesting that striatal activity precedes the execution of well-learned chunks, possibly acting as a gating mechanism for sequential action (Graybiel, 1998; Boyd et al., 2009).

The debate over whether chunking represents a motoric phenomenon or an abstract cognitive strategy underscores the complexity of motor sequence learning. It highlights the interplay between the cognitive processes that organise and recall the order of actions and the sensorimotor processes that execute these actions efficiently. This dichotomy reflects the broader components of motor learning: the spatial-ordinal aspect and the sensorimotor aspect. Understanding how chunking facilitates the integration of these components is essential for a comprehensive view of motor skill acquisition (Ariani and Diedrichsen, 2019; Zimnik and Churchland, 2021). In the following section, I explore these components of learning in detail.

1.3.2 Components of Learning

The process of motor learning can be broadly categorized into two fundamental components: the spatial-ordinal or "what-to-do" component, and the sensorimotor or "how-to-do" component. The spatial-ordinal aspect refers to the learner's ability to understand and sequence the actions required for task completion. This component emphasizes the order and spatial arrangement of movements within a motor sequence, such as the steps involved in playing a musical scale or executing a complex athletic manoeuvrer (Karni et al., 1995; Abrahamse et al., 2010). It is deeply rooted in cognitive processes, where the motor system learns to map external goals

onto specific motor actions, and heavily relies on both working memory and attention. On the other hand, the sensorimotor component pertains to the fine-tuning of the physical execution of these movements. It focuses on performance quality, including timing, force, and fluidity, and is shaped by repeated practice and feedback, allowing efficient and adaptive execution of learned actions (Schmidt, 1975; Fitts, 1964). Together, these two components underpin the holistic development of motor skill, where both the knowledge of what actions to perform and the skilful execution of those actions are mastered over time.

The Spatial-Ordinal Component of motor sequence learning includes the learner's capacity to understand and reproduce the correct spatial arrangement and temporal order of movements. This aspect is vital for both simple and complex tasks, whether it's typing on a keyboard or playing a musical instrument. The ability to map movements spatially ensures that the sequence can be accurately repeated. Research has demonstrated that humans, when learning motor sequences, tend to prioritize understanding the spatial order of movements before perfecting their execution, as this provides the foundation for later refinement of timing and precision (Verwey, 2001; Abrahamse et al., 2013). Furthermore, spatial understanding is essential for motor tasks that require precision, such as surgical procedures or sports, where any deviation in spatial order can lead to errors or suboptimal performance.

Spatial learning becomes especially important in tasks where movements are directed toward specific spatial targets. For instance, in tasks that involve reaching or saccadic eye movements, the accuracy with which spatial positions are remembered and executed is a determinant of success. The precision of spatial-ordinal learning is thought to be supported by the brain's ability to map both egocentric (body-centered) and allocentric (external environment-centered) spatial frameworks. This dual spatial representation allows for more flexible motor control and is necessary for tasks that involve the transformation of spatial representations from one context to another, such as when an individual is required to perform learned movements across environments (Jeannerod, 1988; Kalaska and Crammond, 1992).

In empirical studies of motor sequence learning, spatial-ordinal information

has been shown to drive early improvements in sequence execution. As learners become familiar with the order of the required movements, performance becomes more consistent, and the motor system begins to encode these spatial patterns more robustly. Importantly, learning the spatial-ordinal structure of a task also facilitates transfer to new tasks. For instance, once a spatial sequence is learned in one hand, it can often be transferred to the other hand, demonstrating the generalisation capabilities of the motor system (Hardwick et al., 2013). This highlights the critical role of spatial-ordinal learning in forming flexible motor representations.

Overall, spatial-ordinal learning forms the basis of motor sequence learning, upon which more complex, sensorimotor refinements are built. The process of encoding these spatial relationships enables the learner to transition from conscious, effortful execution to more automatic, fluid movements as the sequence becomes ingrained in both memory and motor control systems.

The Sensorimotor Component of motor sequence learning involves the integration of sensory input with motor output, which is critical for refining movement execution over time. While spatial-ordinal learning focuses on the order and arrangement of movements, the sensorimotor aspect emphasizes the quality of execution - timing, force, and coordination. Sensorimotor integration ensures that movements are adapted in real-time based on feedback from the body and environment. This feedback loop allows for ongoing adjustments that transform the task-relevant kinematics (speed, smoothness, etc.), which is essential for developing highly skilled behaviours (Adams, 1971; Shadmehr and Krakauer, 2010).

Sensorimotor learning is particularly crucial for tasks that require fine motor control, such as playing a musical instrument or performing delicate surgical manoeuvrers. These tasks demand precise timing and coordination, which can only be achieved through the close integration of sensory information (such as proprioception and vision) with motor commands. The role of feedback in this process cannot be understated - real-time corrections allow for adjustments that enhance accuracy and performance (Wolpert and Flanagan, 2001). As learners practice, they become more adept at using this feedback to make micro-corrections that optimize

movement, leading to improved efficiency and reduced variability in execution.

Evidence suggests that sensorimotor learning plays a critical role in the later stages of motor sequence learning, as performance becomes increasingly dependent on the fine-tuning of movements. During practice, the brain refines its internal models, which predict the outcomes of actions and adjust motor commands accordingly. These internal models are crucial for the smooth and efficient execution of complex motor sequences, as they allow anticipatory adjustments rather than reactive corrections (Flanagan and Wing, 1997; Wolpert and Flanagan, 2001). In this sense, the sensorimotor system is continuously learning how to better predict and control movement outcomes based on past experience.

In sequence learning tasks, such as the serial reaction time task (SRTT), participants gradually improve their performance as they refine the coordination between sensory input and motor execution. With practice, the motor system becomes more efficient at integrating sensory feedback, resulting in faster and more accurate responses. This integration of sensory feedback not only facilitates the learning of specific motor tasks but also allows for the generalisation of learned skills to new contexts, demonstrating the flexibility of the sensorimotor system.

1.3.3 Stages of Sequence Learning

A growing body of evidence suggests that motor sequence learning can be divided into three distinct stages, as first proposed by Fitts and Posner in 1967. These stages, typically labelled as the cognitive, associative, and autonomous stages, reflect the progression from effortful, conscious control of movements to automatic execution. Although different theoretical models use varying terminology, the core idea of a phased learning process remains consistent across studies.

The first stage, often referred to as the cognitive, fast or early stage, is marked by rapid improvements in performance as the learner begins to understand the basic structure of the motor sequence. During this phase, performance gains are typically seen in the form of reduced reaction times and increased accuracy in executing the sequence (Doyon and Ungerleider, 2002; Karni et al., 1995). These improvements are largely driven by cognitive processes, as learners rely heavily on explicit mem-

ory to remember the sequence and consciously guide their movements.

During the early phase, participants tend to make frequent errors as they attempt to memorize the order of actions. However, with repeated practice, the frequency of these errors decreases, and performance becomes more consistent (Karni et al., 1995). This rapid improvement in the early stage is thought to be due to the formation of initial motor representations, which serve as the foundation for further refinement in the later stages of learning.

The intermediate or associative stage follows the initial cognitive phase and marks a period of consolidation and refinement. In this stage, learners begin to make fewer errors and rely less on conscious control, as their movements become more fluid and consistent. Attention is gradually redirected from remembering the order of actions to improving timing and coordination. Performance becomes more stable, and learners start to recognize and correct their own mistakes, reflecting a shift from explicit control to more efficient sensorimotor integration (Fitts and Posner, 1967; Doyon and Ungerleider, 2002). This phase is often associated with the strengthening of cortico-striatal pathways and increased involvement of the sensorimotor cortex as skill becomes more embodied.

The late (slow) learning stage is characterized by the automation of motor sequences. This stage, also referred to as the autonomous phase, is characterized by subtle improvements to performance, e.g. gains in efficiency, movement fluidity, and precision, rather than large reductions in response times. The motor sequence, now more deeply encoded in neural circuits, operates with minimal conscious oversight, relying predominantly on implicit memory systems. This shift reflects the gradual internalization of motor commands, particularly as neural activity becomes focused in brain regions associated with habit formation and automated actions, such as the basal ganglia and the cerebellum (Doyon and Ungerleider, 2002; Karni et al., 1995). The refinement of internal models within these regions supports smoother execution, even under complex and variable task conditions.

In the late stage, the motor system demonstrates greater adaptability to perturbations, further refining the timing, coordination, and overall quality of movements.

However, this phase is not solely about the automation of movement. Evidence suggests that the representation of motor sequences may shift from cortical areas involved in planning, such as the premotor cortex (PMC), toward more execution-based regions, including the primary motor cortex (M1) (Hikosaka et al., 2002; Doyon and Ungerleider, 2002). The PMC, responsible for integrating sensorimotor information, continues to play a role in fine-tuning motor sequences, while M1 is engaged in executing actions (Hardwick et al., 2013).

The point of motor learning plateau, often referred to as the retention stage, is characterized by nearly automatic execution of motor sequences. How skilled behaviour transitions from goal-directed to habitual activity has been extensively studied (Adams, 1971; Dickinson and Balleine, 1994). During the early stages of motor learning, actions are selected based on expected outcomes, meaning that individuals choose movements they believe will lead to success. This goal-directed system is computationally intensive and relies on explicit memory processes. As learning progresses, however, skill becomes habitual, based on past successes rather than immediate outcomes (Daw et al., 2005; Balleine and Dickinson, 1998).

Neural correlates of the habitual and goal-directed systems involve the basal ganglia and prefrontal cortex, respectively (Balleine and Dickinson, 1998). When participants are under time pressure and forced to make rapid decisions, they are more likely to rely on habitual responses. In contrast, when given more time, participants revert to goal-directed decision-making, which can result in better task performance (Balleine and Dickinson, 1998). These findings suggest that habitual and goal-directed systems operate in parallel, each utilised based on task demands. This co-utilisation allows for continued refinement of motor skills and maintenance of performance even after the early and late learning phases have concluded.

Long-term retention and maintenance of motor sequence skills requires ongoing engagement with the learned material to prevent skill decay. Long-term maintenance involves processes that ensure motor skills remain accessible and efficient over extended periods. Regular rehearsal and reactivation of motor memories are essential to counteract forgetting and reinforce the neural representations responsi-

ble for sequence execution (Censor et al., 2014; Brunton et al., 2009).

A critical aspect of long-term maintenance is the concept of reconsolidation, where previously consolidated motor memories are retrieved and updated based on new experiences or environmental changes. Reconsolidation allows the brain to refine and adjust motor sequences, ensuring that they remain adaptable and relevant to the individual's current motor demands (Nader et al., 2000; Gelman, 2003). For instance, in sports or rehabilitation, periodic revisiting of motor tasks allows individuals to not only maintain proficiency but also improve upon previously learned skills by incorporating novel feedback into their motor programs.

Another factor influencing long-term maintenance is the spacing of practice sessions. Distributed practice, where learning is spread out over time, has been shown to be more effective for long-term retention than massed practice, where learning occurs in concentrated bursts (Karni et al., 1995; Baddeley and Longman, 1978). The spacing effect allows for the gradual strengthening of motor memories, reducing the likelihood of forgetting and promoting the durability of the learned sequences. This principle is particularly relevant in skill-based domains such as surgery, where proficiency must be maintained over long periods without regular engagement with the specific motor task.

1.4 Mechanisms of Motor Sequence Consolidation

Motor sequence learning stage progression is supported by the consolidation of movement patterns into stable motor memories. Consolidation refers to the processes that stabilise and enhance motor memories after initial learning, making them more resistant to interference and decay. In the context of motor sequence learning, consolidation can occur during active practice (online), during periods of wakeful rest (offline), and during sleep. Each of these phases contributes uniquely to the development of skilled performance. This section examines the mechanisms underlying motor sequence consolidation, focusing on behavioural and neurological evidence for online consolidation during practice, offline consolidation during wakeful rest, and overnight sleep consolidation.

1.4.1 Online Consolidation During Practice

Online consolidation in motor sequence learning refers to the rapid improvements in performance observed during active practice. Behavioural studies have consistently shown that participants exhibit significant reductions in reaction time (RT) and improvements in accuracy while practising motor sequences within a single session (Karni et al., 1995; Dayan and Cohen, 2011). Improvements often follow a characteristic learning curve, with rapid gains early in practice that gradually plateau - a phenomenon reflecting the initial encoding and refinement of motor sequences.

For example, in serial reaction time tasks (SRTT), participants respond to sequential stimuli by performing corresponding motor actions. As practice progresses, RTs decrease, indicating that the sequence is becoming more ingrained and that movements are being executed more efficiently (Nissen and Bullemer, 1987). Similarly, in discrete sequence production (DSP) tasks, participants show faster execution of practised sequences, suggesting that they are consolidating the motor patterns required for the sequence during practice (Verwey, 2001).

Neural mechanisms underlying online consolidation during motor sequence practice involve a widespread network including the prefrontal cortex, premotor cortex, parietal cortex, basal ganglia, and cerebellum (Hikosaka et al., 2002; Doyon and Ungerleider, 2002). These areas are associated with attention, planning, and the execution of movements. As practice continues and performance improves, there is a shift in neural activation patterns. Activity becomes more focused in the primary motor cortex (M1) and supplementary motor area (SMA), regions directly involved in motor execution and sequence representation . This neural reorganization reflects the consolidation of motor sequences into more efficient neural circuits, enabling faster and more automatic execution of the learned sequences (Karni et al., 1995).

At the cellular level, synaptic plasticity mechanisms such as long-term potentiation (LTP) are believed to contribute to the strengthening of synaptic connections within motor circuits during practice (Rioult-Pedotti et al., 2000). These changes enhance the responsiveness of motor neurons to specific patterns of input, facilitating the consolidation of motor sequences during online learning.

1.4.2 Offline Consolidation During Wakeful Rest

Offline consolidation during wakeful rest refers to performance improvements in motor sequence learning that occur between practice sessions without additional practice or sleep. Behavioural studies have demonstrated that brief rest periods interleaved within practice sessions can lead to significant gains in performance known as micro-offline gains (Antonenko et al., 2013; Bonstrup et al., 2019). Bönstrup et al. (2019) found that participants practising a finger-tapping motor sequence exhibited marked improvements in speed and accuracy after short rest intervals compared to continuous practice without breaks. Improvements suggest that the motor system continues to consolidate and optimize the learned sequences during wakeful rest, enhancing performance upon resumption of practice. Moreover, the magnitude of offline gains during wakeful rest is pronounced in the early stages of motor sequence learning, where initial encoding is followed by rapid consolidation processes (Hotermans et al., 2006). This emphasizes the importance of rest periods in facilitating the consolidation of newly acquired motor sequences.

Neurological evidence supports the role of wakeful rest in the consolidation of motor sequences. During rest periods following motor sequence practice, there is continued activation in motor-related brain regions, including M1, SMA, and the basal ganglia (Albouy et al., 2008). This ongoing neural activity suggests that the brain is replaying or reinforcing the motor sequences learned during practice. Electroencephalography (EEG) studies have identified patterns of oscillations during wakeful rest that correlate with offline performance gains. For example, increased beta-band oscillatory activity in motor regions during rest has been associated with improved motor sequence performance, indicating that neural synchronization may play a role in consolidating motor memories (Penhune and Steele, 2012). Transcranial magnetic stimulation (TMS) research shows changes in cortico-spinal excitability during wakeful rest following motor sequence learning (Robertson, 2005), which reflect synaptic modifications in motor pathways that contribute to the stabilisation and enhancement of motor sequences during offline periods.

1.4.3 Overnight Sleep Consolidation

Overnight sleep has been extensively studied for its role in consolidating motor sequences. Behavioural experiments have shown that participants often exhibit significant performance improvements in motor sequence tasks after a night of sleep compared to equivalent periods of daytime wakefulness (Walker et al., 2002; Fischer et al., 2002). These improvements include faster execution speeds, increased accuracy, and reduced variability in movement timing.

For example, Walker et al. (2002) demonstrated that participants practising a finger-tapping sequence showed substantial gains in speed and accuracy after sleep but not after an equivalent period of wakefulness. In tasks such as the serial reaction time task (SRTT), participants not only maintain performance after sleep but often display improvements without additional practice, reflecting the strengthening of motor memories during sleep (Schapiro et al., 2018).

Further, the amount of performance improvement correlates with specific sleep stages. Studies have found that both non-rapid eye movement (NREM) sleep, particularly stage 2 sleep spindles, and rapid eye movement (REM) sleep contribute to motor sequence consolidation (Fischer et al., 2005; Nishida and Walker, 2007). Enhanced sleep quality, particularly with an emphasis on increasing slow-wave sleep (SWS), can lead to more pronounced gains in motor sequence learning, suggesting a direct relationship between sleep architecture and learning efficiency.

The influence of sleep on motor sequence learning is not uniform and can be affected by factors such as sleep quality, the complexity of the motor task, and individual differences in sleep architecture. Disruptions in sleep, such as sleep deprivation or fragmentation, have been shown to impair the consolidation of motor skills, leading to poorer performance outcomes the following day (Fischer et al., 2002). Conversely, improving sleep quality can enhance learning efficiency, indicating the importance of optimal sleep conditions for motor memory consolidation.

Sleep-dependent consolidation of motor sequences involves complex neural processes that reorganize and strengthen motor memory representations. Neuroimaging studies using fMRI have shown that after sleep, there is increased ac-

tivation and functional connectivity in motor regions associated with the learned sequences, such as the primary motor cortex (M1), supplementary motor area (SMA), and the cerebellum (Tamaki et al., 2013). Sleep promotes synaptic plasticity and neural reactivation of motor sequences. During NREM sleep, particularly slow-wave sleep (SWS), replay of neural patterns associated with the practised sequences occurs in motor-related areas and the hippocampus, reinforcing synaptic connections (Walker and Stickgold, 2005; Born and Wilhelm, 2010; Diekelmann and Born, 2010). This replay enhances synaptic plasticity, making the learned motor sequences more robust and integrated into long-term memory networks. Sleep spindles, characteristic of stage 2 sleep, have been linked to synaptic consolidation processes that strengthen motor memories (Barakat et al., 2013).

Additionally, REM sleep is thought to contribute to the integration and refinement of motor sequences into existing memory networks, facilitating the generalization and transfer of skills (Stickgold and Walker, 2007; Smith, 2001). REM sleep may facilitate synaptic pruning and the optimization of neural networks, enhancing motor efficiency and allowing the motor system to execute learned sequences with greater automaticity and precision. This balance between SWS and REM sleep stages ensures that motor memories are not only retained but also refined for future performance (Walker and Stickgold, 2005).

Table 1.2: Overview of consolidation mechanisms and associated neural regions

Mechanism	Description	Neural Correlates
Online	Consolidation occurring during active practice; improvements emerge during repetition.	Motor cortex (M1), Striatum (Dayan and Cohen, 2011)
Micro-offline	Rapid consolidation during short rest periods between practice bouts, even within a single session.	Supplementary Motor Area (SMA), Striatum (Bonstrup et al., 2019)
Sleep-dependent	Consolidation across longer timescales (e.g., overnight), with delayed performance gains.	Hippocampus, M1 (Walker and Stickgold, 2003, 2005)

1.4.4 How We Move On: Motor learning to motor control

Consolidation over online, offline, and sleep-dependent phases stabilises motor sequences and refines performance in unique ways. Online consolidation strengthens initial motor representations, offline consolidation during wakeful rest allows for rapid enhancements and stabilisation, and overnight sleep facilitates long-term retention and integration of motor sequences into broader motor networks.

While consolidation processes stabilise and enhance motor sequences, effective execution of these sequences relies on motor control. In the next section, I explore the relationship between motor control and motor learning, examining how neural and cognitive mechanisms underlying movement execution interact with learning processes to produce coordinated and efficient motor behaviours.

1.5 Motor Control, Redundancy and Noise

Motor control involves the cortico-spinal coordination of muscles and joints to execute actions. A central challenge in motor control is redundancy, where the body's numerous degrees of freedom provide myriad ways to perform the same movement (Bernstein, 1947a; Latash, 2010). For instance, a simple reach can be executed using many different muscle combinations and joint angles. The brain must select the most appropriate strategy for a given task from the large number of possibilities at hand, often gravitating toward stereotypical movement patterns, such as the slight curvature observed in reaching movements (Morasso, 1981).

A second challenge is motor noise, which refers to variability in the motor system that increases with the magnitude of movement (Harris and Wolpert, 1998). This phenomenon, known as signal-dependent noise, arises from fluctuations in neural signals, muscle recruitment variability, and biomechanical factors. For example, when reaching to lift a full cup versus an empty one, the heavier cup requires more force, which in turn introduces more variability into the movement. The motor system must compensate for this increased noise to maintain precision, especially during fast or forceful actions.

Motor sequence learning is a crucial aspect of managing these control chal-

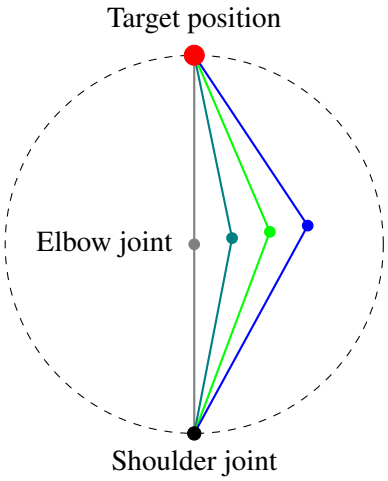

Figure 1.1: Redundant joint angle configurations reaching a common target. The same hand position can be achieved with multiple elbow angles, demonstrating the problem of redundancy in human motor control.

lenges. When a sequence of actions is required - such as typing a password or playing a melody on a piano - the brain must not only plan each individual action but also ensure the fluid execution of the entire sequence. Initially, movements within a sequence are executed slowly and deliberately, requiring attention and conscious control (Karni et al., 1995; Hikosaka et al., 2002). As practice continues, these movements become automated and shift from being controlled by higher-level cortical regions to being embedded in subcortical circuits, such as the basal ganglia and cerebellum (Doyon and Ungerleider, 2002; Karni et al., 1995). To do this, motor sequence learning must navigate redundancy and noise by forming reliable internal models of movement dynamics. These models allow the brain to predict the sensory consequences of motor commands and adjust movements accordingly, helping to overcome temporal delays in sensory feedback (Kawato, 1999).

In short, motor control is not just about selecting and executing individual actions but also about learning and optimizing sequences of movements. By mastering sequences through practice, the motor system develops internal models that enable efficient, reliable performance across a range of tasks. This sets the stage for the more complex challenge of coordinating multiple limbs and effectors, which I explore in the next section.

1.6 Multi-Effector Motor Control

Motor control complexity scales as we coordinate multiple effectors, such as bimanual tasks or walking and adjusting arm movements simultaneously. Multi-effector control refers to the coordination of multiple limbs or digits at once, ensuring that each move effectively to achieve a shared goal. This process requires the integration of sensory input and motor commands across various regions of the brain, e.g. the premotor cortex and supplementary motor area plan and initiate movements in bimanual tasks, while the cerebellum tunes coordination to prevent errors and facilitate smooth execution of movements (Koziol et al., 2014).

Understanding how the motor system controls multiple effectors - whether coordinating limbs or muscles across joints - requires sophisticated models that explain the underlying neural computations. Two major frameworks, optimal control theory and the internal model framework, have been pivotal in providing theoretical explanations for how the brain manages the complexities of multi-effector control.

Optimal Control Theory proposes that the brain continuously seeks to minimize a cost function during movement execution. This cost function encompasses factors such as effort, noise, and accuracy, balancing them to achieve an efficient and adaptive movement outcome (Todorov and Jordan, 2002; Scott, 2004). In multi-effector tasks, where coordination between limbs is crucial, the motor system must account for both internal constraints (e.g., muscle strength, limb length, and redundancy in degrees of freedom) and external environmental perturbations (e.g., gravity or friction). Rather than aiming for flawless precision, the motor system optimizes performance by finding the best compromise between competing demands.

For example, consider a person reaching for an object with one hand while using the other to stabilise themselves. Optimal control theory predicts that the brain will allocate control resources based on task relevance: the reaching arm, which is central to the task, receives higher precision and tighter control, while the stabilising arm may tolerate greater variability. This strategic distribution minimizes overall effort while ensuring successful task performance (Todorov and Jordan, 2002).

Extensions of optimal control theory incorporate both feedback and feedforward mechanisms that govern multi-effector coordination. In feedback control, sensory signals are used to correct ongoing movements. In feedforward control, the brain generates predictions about the consequences of motor commands and adjusts them in advance of sensory feedback. By integrating both mechanisms, the motor system can anticipate and correct errors dynamically, allowing for smooth coordination across multiple effectors. This interplay is especially important in complex tasks such as playing the piano or typing, where multiple limbs or digits must act in synchrony toward a shared objective (Franklin and Wolpert, 2008).

The Internal Model Framework emphasizes the brain's ability to predict the outcomes of its motor commands. According to this framework, the brain builds and updates internal representations (models) of the body and the external environment (Kawato, 1999; Wolpert and Flanagan, 2001). Internal models serve two key purposes: forward models predict the sensory consequences of motor actions, and inverse models determine the motor commands needed to achieve a desired outcome.

In multi-effector tasks, internal models allow the brain to coordinate movements across different limbs by predicting how changes in one effector will impact the others. For example, when reaching with one hand while balancing on one leg, the brain uses internal models to predict how shifts in the body's centre of mass will affect stability and adjust accordingly (Kawato, 1999; Shadmehr and Krakauer, 2010). This predictive capability is particularly important in tasks that involve rapid, synchronized movements, such as playing a musical instrument or performing a gymnastics routine. Here, the timing and coordination between limbs must be finely tuned, and the internal models allow the motor system to adjust movements even before feedback from the limbs is available.

Internal models also contribute to the brain's ability to adapt to new conditions. In situations where the dynamics of the environment change - such as when using tools or performing tasks in altered gravity - the brain can update its internal models to reflect the new context, ensuring accurate multi-effector coordination. This adaptability is crucial in real-world scenarios, where changes in the environment

require rapid recalibration of motor commands.

Internal models provide a theoretical basis for understanding how the brain manages the complex task of controlling multiple limbs while maintaining skill. By simplifying control through predictive models and a mechanism that I will explore in the next section, motor synergies, the motor system is able to achieve smooth, coordinated movement across different effectors.

1.7 Motor Synergy

What is synergy? Synergy is a term that frequently appears in the field of motor control, but it often lacks a clear definition. Many papers claim to investigate features or mechanisms of synergy without fully explaining what it entails. In some instances, synergy is loosely described as "coordinated action," though coordination itself is rarely defined. At other times, synergy is operationally defined as variables that "change together," where this co-variation can range from simple linear scaling to more complex indices of coherence. The concept has been discussed at length in various papers, chapters, and books (Bernstein, 1947b; Latash et al., 2007; Turvey, 2007; Latash, 2008a, 2020; Tresch and Jarc, 2009; Bruton and O'Dwyer, 2018). Despite these discussions, consensus on its meaning remains elusive.

Synergy has been a foundational concept in studies of biological action for centuries. Notably, St. Gregory Palama, an Orthodox Christian theologian from the 13th and 14th centuries, used the term to describe the cooperative action of God and Man in human salvation (Palamas, 1983, 1988; Meyendorff, 1974). He described a negative co-variation between the efforts of God and Man, a concept surprisingly close to modern interpretations of synergy in motor control.

By the 19th century, terms like synergy, asynergia, and dyssynergia were used in clinical neurology to describe coordination impairments in patients with neurological disorders (Hughlings Jackson, 1889; Babinski, 1899). For example, Babinski associated synergies with normal cerebellar function, an insight supported by relatively modern research (Thach et al., 1992; Welsh and Llinas, 1997; Ting and McKay, 2008). However, during the mid-20th century, the meaning of synergy in

clinical contexts shifted to describe poor coordination patterns, particularly those seen in stroke patients (Bobath, 1978; Dewald et al., 1995). These "pathological synergies" included flexion and extension synergies that were seen as impediments to functional movement, leading to a use of the term as a negative.

Nikolai Bernstein laid the groundwork to understand synergy in the context of motor control, describing it within a multi-level hierarchical framework. He associated synergies with the organization of muscles and the maintenance of dynamic stability during movement. His insights into the "Level of Synergies" highlighted how the central nervous system copes with the inherent unpredictability of both internal and external forces acting on the body (Bernstein, 1947b, 1967).

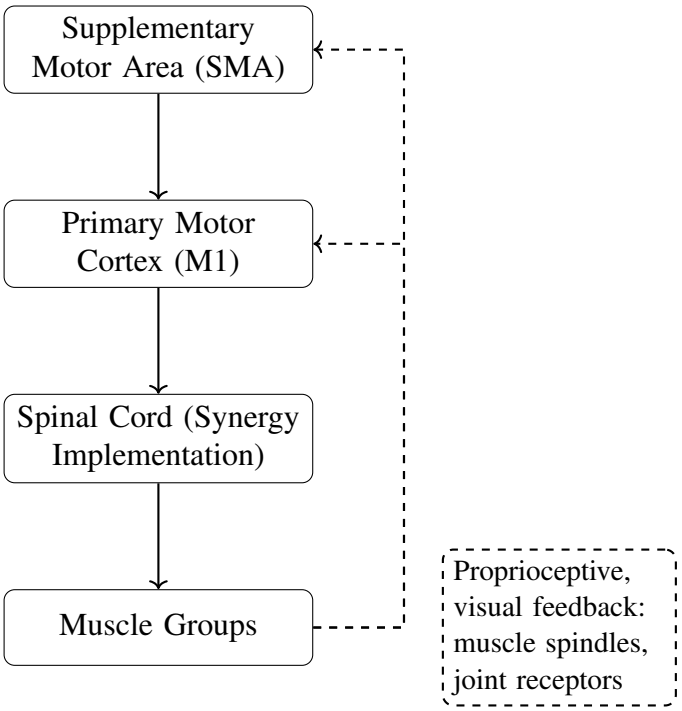

Figure 1.2: Model of synergy pathways from cortex to muscle, with feedback.

An alternative perspective, deeply rooted in engineering and control theory, posits that the brain generates "internal models" to plan, predict, and execute motor commands at the level of muscles and joints (Wolpert et al., 1995; Kawato, 1999; Shadmehr and Wise, 2005). Although Bernstein's approach and control theory are based on fundamentally different assumptions, they both contribute valuable insights into the mechanisms underlying motor control. In this thesis, motor control

is examined as a critical aspect of motor sequence learning, with motor synergy positioned as a specialized control mechanism that enhances motor skill performance. Specifically, synergy is conceptualized here as the coordination and co-variation of multiple kinematic features across the hand during skilled movement. This definition reflects the dynamic interactions that explain gains in motor performance and invites further comparison with broader theories of motor control, including optimal control and internal models.

An important distinction in the literature is between synergy as a descriptive construct and synergy as an explanatory control mechanism (Tresch et al., 1999). In the descriptive sense, synergies are observed as systematic co-variation between effectors - such as fingers or limbs - across repeated movements. This co-variation may arise passively from biomechanical constraints or statistical regularities without implying active neural coordination. In contrast, the explanatory view posits that synergies reflect centralized neural strategies, wherein the motor system actively groups effectors into functional units to reduce computational complexity. This thesis adopts the latter perspective, treating motor synergy as a purposeful control structure that enhances performance in skilled motor sequences.

This explanatory perspective aligns with, but also challenges, aspects of optimal control theory and internal model frameworks. While optimal control theory explains motor behaviour as the outcome of minimizing task-relevant costs (e.g., effort, noise, inaccuracy), it typically assumes flexible, continuous control over individual effectors. In contrast, a synergy-based model may impose modular constraints on effector coordination, potentially limiting flexibility in favour of stability and efficiency. Whether synergies emerge from optimal control principles or represent a distinct layer of motor organization remains debated. Nonetheless, both frameworks converge on the idea that the brain achieves coordinated behaviour through economical, adaptive strategies, especially in multi-effector tasks.

1.7.1 Synergy Expression

Synergies are expressed across various domains - kinematic, muscular, cognitive, and neurological - offering a rich framework for understanding how humans pro-

duce coordinated movements. The observation of synergies across these dimensions dates back to early explorations in motor control, where researchers began to recognize that movement is not the result of isolated muscle activity but rather the product of organized patterns across muscle groups, joints, and neural networks.

Historically, synergies were first studied in kinematic terms, focusing on how the body achieves fluid, coordinated movement across joints and limbs. As technology advanced, muscle activity was increasingly understood through electromyography (EMG), allowing for direct observation of how different muscle groups are activated in concert during complex actions. These muscular synergies revealed that groups of muscles are recruited together as functional units, providing efficiency and stability in movement.

Simultaneously, studies of the brain's role in motor control have revealed that synergies are not merely a byproduct of physical constraints but are actively orchestrated by the CNS. Neurological research has shown that synergies are encoded in the motor cortex and other regions of the brain, where preconfigured patterns allow for rapid and flexible motor responses. Cognitive research, too, has added to this understanding, highlighting how synergies are involved not just in motor execution but also in planning, decision-making, and even skill acquisition, reinforcing the multidimensional nature of synergy expression. I will touch on each of these modes of expression in the sections below.

Kinematic indicators of synergy involve the spatial and temporal coordination of movement patterns and reveal how different segments of the body move in unison to accomplish a specific task. Kinematics refers to the study of motion without considering forces, focusing on aspects like velocities and trajectories.

A classic example of kinematic synergy can be seen in reaching movements. When a person reaches for an object, the shoulder, elbow, and wrist move in a coordinated fashion, minimizing the degrees of freedom and simplifying the control problem. Studies have shown that even in multi-joint movements, individuals tend to use highly stereotyped movement trajectories that reflect underlying motor synergies (Morasso, 1981; Bernstein, 1967). Additionally, in locomotion, kinematic

synergies are observed in the coordination of limbs to maintain balance and stability during gait, with the legs and arms moving in a synchronized fashion to optimize energy expenditure (Winter, 1990).

Kinematic synergies are also apparent in skilled tasks, such as playing a musical instrument or typing. In such tasks, finger movements are often coordinated to form a sequence of actions that minimize errors and increase speed (Latash, 2008a). The ability to optimise joint angles and trajectories while maintaining task accuracy is a hallmark of kinematic synergy expression, highlighting the motor system's capacity to streamline control across multiple degrees of freedom.

Muscular indicators of motor synergy refer to the coordinated activation of multiple muscles to produce effective movement. Instead of controlling each muscle individually, the CNS groups muscles into synergies, allowing more straightforward and robust movement execution. These synergies are often flexible, meaning that the same group of muscles can be used to produce different movements depending on the task demands.

Electromyography (EMG) studies provide insight into how muscle synergies are expressed. EMG recordings show that certain muscle groups are activated together in a highly correlated manner during specific movements. For instance, in locomotion, the muscles of the lower limbs exhibit coordinated patterns of activation to produce smooth and rhythmic walking motions (Overduin et al., 2008, 2012; Bizzi et al., 2008). Even in tasks that require fine motor control, such as grasping, muscle synergies can be observed, with specific hand muscles activating together to stabilise and manipulate the object (Delis et al., 2018).

Cognitive and neurological indicators of motor synergy involve physiological markers of higher-order motor planning and decision-making in generating coordinated movements. The planning of actions, particularly those involving multi-effector control, relies heavily on working memory, attention, and decision-making processes. For instance, studies have shown that the ability to recall sequences and apply them in a motor context is highly dependent on cognitive load and the involvement of prefrontal cortical areas (Yokoyama et al., 2019). These areas of the

brain contribute to the selection and organisation of motor synergies, ensuring that groups of muscles are activated in an efficient and task-relevant manner. Moreover, as movements become more automatic with practice, the cognitive load decreases, reflecting the shift from active decision-making to more ingrained, habitual responses. This cognitive shift is associated with a reduction in the involvement of prefrontal regions, as control over movements transitions to subcortical structures, such as the basal ganglia which are theorised to be involved in the selection and initiation of motor synergies, especially in the context of sequential tasks (Graybiel, 2008; Shadmehr and Krakauer, 2010).

Further, the cerebellum contributes to the fine-tuning of motor synergies, ensuring that movements are smooth and well-coordinated. It does so by constantly integrating sensory feedback and adjusting ongoing movements to correct for errors (Manto et al., 2012). The cerebellum is especially important for timing and coordination in tasks that require precise synchronization of multiple limbs or muscle groups. Disruptions in cerebellar function often lead to a breakdown in motor synergies, resulting in uncoordinated or jerky movements, as seen in disorders such as ataxia (Wolpert and Flanagan, 2001).

The premotor cortex (PMC) and supplementary motor area (SMA) are also heavily involved in the planning and execution of synergies. The PMC integrates sensory and motor information to organize movements across multiple effectors, while the SMA is particularly important for coordinating internally generated sequences of actions, such as those required for typing or playing an instrument (Hikosaka et al., 2002; Doyon and Ungerleider, 2002). The interaction between these cortical regions and subcortical structures allows the motor system to dynamically adjust synergies based on both internal goals and external feedback.

1.7.2 Synergy Quantification

To comprehensively appreciate the expression of synergies, it is crucial to quantify them accurately across various domains. Quantifying synergies offers insights into the underlying physiological and mechanical strategies that govern coordinated motor behaviours. Several advanced techniques have been developed to achieve this,

ranging from mathematical factorization methods to models rooted in non-linear dynamics and dimensionality reduction. This section delves into these methods – Non-negative Matrix Factorization (NNMF), the Uncontrolled Manifold (UCM) hypothesis, and non-linear dimensionality reduction techniques – all of which have become critical for understanding the structure of motor synergies.

Principal Component Analysis (PCA) is a widely-used method in detecting and analysing motor and muscle synergies, especially when processing large datasets from electromyography (EMG) recordings. PCA identifies the principal components, or key modes of variance, within a dataset, enabling researchers to reduce the dimensionality of EMG data without losing significant information. This reduction is particularly beneficial in motor control studies, where multiple muscles contribute to complex motor tasks. By using PCA, muscle synergies can be extracted as linear combinations of muscles contributing to the same motor function, allowing an understanding of how muscle groups are activated together during movement (Cheung et al., 2005; Castellini and van der Smagt, 2013).

PCA is particularly valuable in synergy analysis because it identifies patterns of co-activation across muscles, which can be interpreted as muscle synergies. These synergies are defined as the coordinated activation of muscles to achieve a specific motor goal. PCA elucidates these patterns by reducing the complexity of EMG signals into a set of orthogonal components. In studies of locomotion or grasping, for example, PCA has revealed that a small number of synergies can account for most of the variance in muscle activity, suggesting that the nervous system simplifies control by activating muscles in groups rather than individually (Cheung et al., 2005; Castellini and van der Smagt, 2013). Such findings indicate that PCA is crucial for isolating key muscle patterns that contribute to effective motor control in both simple and complex tasks.

PCA's ability to represent the majority of a dataset's variance in a reduced number of components is highly advantageous in motor control research. It allows researchers to focus on the critical aspects of motor behaviour. For instance, when analysing gait or postural adjustments, PCA can distinguish between healthy and

pathological movement patterns and identify the primary muscles responsible for maintaining balance or producing movement (Berman et al., 2014; Cappellini et al., 2010). Moreover, by extracting synergies through PCA, researchers can compare patterns across different conditions or patient groups, providing insights into the neurological basis of movement disorders.

Despite the many advantages of PCA in detecting synergy, it has several shortcomings. One of the primary limitations is that PCA assumes linear relationships in the data, meaning it may fail to capture more complex, non-linear synergies that are present in real-world motor tasks. This linearity assumption can overlook important interactions between muscles or motor patterns that do not follow straightforward, linear trends (Sanger, 2000). Additionally, PCA's sensitivity to noise can be problematic, especially when applied to biological data like EMG signals, which tend to have a high degree of variability. Even slight changes in the data can significantly affect the principal components, making it difficult to distinguish genuine motor synergies from noise (Bruton and O'Dwyer, 2018). Another shortcoming is PCA's dependence on the scale of the data. If the input features are not properly normalized, PCA may assign disproportionate weight to features with larger variances, leading to misleading results. This is particularly relevant in muscle synergy studies, where the magnitudes of EMG signals can vary widely across different muscles (Tresch and Jarc, 2009). Finally, while PCA is effective at reducing dimensionality, it is not always interpretable. The principal components themselves may not correspond to physically meaningful muscle synergies, making it challenging to draw direct conclusions about motor control from the results. In cases where more interpretability is needed, alternative methods such as non-negative matrix factorization might be preferred over PCA (Lee and Seung, 1999).

Non-negative Matrix Factorization (NNMF) is a commonly used method in the analysis of muscle activity during motor tasks, particularly through EMG data. NNMF is a powerful factorization technique that decomposes complex multivariate data into lower-dimensional components, which in the context of motor control, can represent underlying muscle synergies. By applying NNMF to EMG recordings, it

is possible to identify the basic building blocks of muscle coordination - synergies - without making a priori assumptions about the data structure (Lee and Seung, 1999; Tresch et al., 1999).

In motor control research, NNMF extracts spatial and temporal patterns of muscle activations that contribute to a specific motor behaviour. For instance, during walking or reaching movements, the technique can break down the complex EMG data into several non-negative components, each corresponding to a distinct muscle synergy (Tresch et al., 2006; Mussa-Ivaldi, 1999). These synergies, once identified, offer a parsimonious explanation of how groups of muscles are activated together to produce smooth, coordinated movements. NNMF's ability to uncover these functional muscle groups has been instrumental in revealing invariant patterns of muscle activity across different motor tasks, which can help understand the modularity of the motor system and how it flexibly adapts to changing task demands (d'Avella et al., 2006; McGowan et al., 2010).

While NNMF is widely appreciated for its interpretability and its biologically plausible constraints (i.e., non-negativity), the challenge remains in its sensitivity to noise and the choice of the number of synergies to retain in a given dataset. This limitation requires careful consideration in practical applications, often demanding cross-validation methods to ensure the robustness of the extracted synergies.

Uncontrolled Manifold (UCM) hypothesis alternatively focuses on the variability in movement performance itself. UCM assumes that the nervous system organizes motor control in such a way that variability in motor outputs is structured: variability that affects task performance is minimized, while variability in dimensions irrelevant to task performance is allowed (Scholz and Schöner, 1999; Latash et al., 2007). Essentially, the UCM framework posits that the nervous system does not control each degree of freedom independently, but rather it controls them in functional groups, allowing for variability that does not affect task success.

Quantifying synergies under the UCM hypothesis involves examining the variability of joint or muscle activation patterns across repeated trials of a motor task. The method computes the variability in the "task space," partitioning it into com-

ponents that do (or do not) affect task performance (Latash, 2010). This allows researchers to assess how motor synergies function in maintaining task-relevant features of movement while allowing flexibility in other motor dimensions. For example, in reaching tasks, the UCM hypothesis has been used to show that joint-level variability is higher in directions that do not interfere with task success (e.g., reaching the target), supporting the idea that the CNS organizes synergies to minimize task-relevant variability (Reisman and Scholz, 2003; Danion et al., 2000).

The strength of the UCM framework lies in its ability to assess motor redundancy not as a limitation but as a feature of efficient motor control, whereby the system organizes itself to exploit variability that does not impact goal achievement. However, one challenge in UCM analysis is its reliance on accurate kinematic and dynamic models of movement to estimate the degrees of freedom involved in a task.

Table 1.3: Comparison of PCA, NNMF, and UCM for Synergy Quantification

Method	Core Principle	Advantages	Limitations
PCA	Identifies orthogonal components that explain maximal variance in EMG or kinematic data.	Data reduction; captures common co-activation patterns; widely used and well understood.	Assumes linearity; sensitive to noise and scale; components may lack interpretability.
NNMF	Factorizes data into non-negative, additive components that reflect muscle synergies.	Intuitive interpretation of synergies; biologically plausible; flexible for different motor tasks.	Sensitive to noise; component number must be chosen carefully; results can vary with initialization.
UCM	Decomposes variability into task-relevant and task-irrelevant subspaces to study control structure.	Directly relates to motor performance; captures structured variability and redundancy.	Requires accurate models of task kinematics; not designed for pattern extraction or factorization.

Non-linear dimensionality reduction techniques are increasingly being applied to study motor synergies, offering a sophisticated method to capture the structure of complex, high-dimensional motor data. Traditional linear methods, such as PCA, often fail to capture the true nature of non-linear relationships within motor data, particularly when examining synergies across different effectors or tasks. Non-linear methods, including techniques like t-SNE (t-distributed stochastic neighbour embedding) and UMAP (Uniform Manifold Approximation and Projection), are ef-

fective in overcoming this limitation by providing a low-dimensional representation of high-dimensional movement data, while preserving its intrinsic structure (van der Maaten and Hinton, 2008; McInnes et al., 2018).

These techniques are particularly useful in exploring how synergies evolve over time or change across different contexts. By applying non-linear dimensionality reduction to multi-effector tasks, researchers can visualise how the motor system organizes synergies dynamically in complex tasks such as bi-manual movements, tool use, or gait (Weaver et al., 2024). These methods also allow for the identification of task-specific synergy patterns that may not be evident using linear techniques, offering insight into the modularity of the motor system.

In this thesis, I apply non-linear dimensionality reduction to understand how synergies restructure in response to motor learning. For example, as individuals practice a new motor task, their synergies may initially be transient or inconsistent, but over time, these patterns become more structured and persistent. Non-linear dimensionality reduction helps capture these transitions (Valero-Cuevas, 2019), providing a visual and quantitative measure of how motor synergies are refined with practice. The use of these techniques has broadened the scope of synergy quantification, enabling the exploration of complex behaviours that would be difficult to capture using more traditional approaches.

1.8 On the Transfer of Motor Synergy Between Skills

The transfer of motor synergies between tasks is an emerging aspect of motor control and learning research. While the broader field of motor learning transfer has been explored in depth, the specific transfer of synergies from one motor skill to another - particularly when tasks are similar but not identical - remains understudied. However, there are several relevant strands of research that highlight the potential mechanisms and contexts in which synergy transfer may occur.

One area of research is the generalisation of motor learning and synergy formation. Generalisation allows motor skills learned in one context to extend to other tasks with similar motor or environmental demands. Although much of the research

in this area has focused on task-level learning, motor synergies, as functional units of motor coordination, may be the underlying mechanism that enables transfer between related skills (Newell, 1991; Sarwary et al., 2015; Schoenfeld et al., 2024). For example, research into motor adaptation suggests that individuals can reuse learned patterns of muscle activation in new tasks with modified demands, potentially transferring synergies as needed (Berniker and Kording, 2008).

Multi-joint movements provide an interesting view into synergy transfer. Tasks such as rowing or swimming, which rely heavily on multi-joint coordination, have shown that synergies developed in one task can be applied to another with similar bio-mechanical demands (Turvey, 1990; Latash, 2008b). This phenomenon has been observed in athletes who transition between related sports, where they can leverage synergies developed for one set of movements to perform in a new context.

As individuals acquire and refine motor skills, the synergies that guide movement become more specialized. Some evidence from motor sequence learning suggests that synergies developed in the early stages of skill acquisition can be applied to more complex sequences, allowing for faster learning and better performance in subsequent tasks (Diedrichsen et al., 2019; Latash, 2008b). This is particularly evident in tasks involving serial elements, such as playing a musical instrument or typing, where synergies established during simpler tasks can scaffold learning in more complex contexts (Verwey et al., 2015; Krakauer and Mazzoni, 2011).

Motor synergies also play an essential role in rehabilitation, particularly for patients recovering from motor impairments such as stroke. In these cases, therapists aim to retrain synergies that can generalise to everyday tasks, allowing patients to regain functional motor abilities. Synergy transfer in this context is critical for the development of adaptive motor patterns that are robust across different activities of daily living (Kleim, 2011; Hong et al., 2021; Zhao et al., 2023).

Research has also shown the flexibility of motor synergies and their adaptability across different tasks in healthy participants. This flexibility is particularly important in sports and musical performance, where high levels of motor coordination are required, and the transfer of motor synergies can significantly impact perfor-

mance (Latash, 2008b; Bernstein, 1967). Understanding how these synergies can be reused or adjusted across tasks provides important insights into skill transference and motor learning.

Despite a growing body of research, studies have only indirectly addressed the transfer of motor synergies. Current research suggests that synergies serve as an adaptive framework for motor skill generalisation, enabling the transfer of learned motor patterns to novel tasks (Kutsuzawa and Hayashibe, 2022; Berg et al., 2023; Herzog et al., 2024). However, further investigation is required to clarify the behavioural and biomechanical mechanisms that govern this transfer.

1.9 Synergy Following a Clinical Event

Movement stability is a common issue in various neurological disorders, most notably in conditions such as ataxia in patients with cerebellar disorders and postural instability in individuals with Parkinson's disease (PD). Normal synergic control involves two critical functions: ensuring the stability of motor actions and reducing this stability in preparation for rapid movement, a process known as anticipatory synergy adjustment (ASA). Neurological patients often demonstrate impairments in both these aspects, manifesting as reduced stability, reflected by a lower synergy index, and diminished agility, characterized by short or limited ASAs (Latash and Huang, 2015). Studies of neurological patients, including those with PD, multiple sclerosis, and multi-system atrophy, have consistently demonstrated impaired stability (Park et al., 2014; Jo et al., 2016b; Ambike et al., 2021). Research on stroke survivors has provided evidence of both impaired and unaltered stability during tasks such as multi-joint reaching and multi-finger force production (Reisman and Scholz, 2003; Jo et al., 2016a; Tomita et al., 2020). These findings underscore the critical role of subcortical circuits in maintaining unimpaired synergies.

ASAs are impaired across all major patient groups, including stroke survivors who display no detectable changes in their synergy index. Reduced ASAs in both duration and magnitude have been implicated as potential contributors to phenomena like freezing, which is characterized by an inability to initiate movement, such

as stepping. This is a significant disabling factor commonly observed in patients with late-stage PD (Falaki et al., 2018, 2023). Research suggests that indices of synergic control are highly sensitive to early stages of neurological disorders, often preceding clinical manifestations. For instance, similar changes in synergy indices and ASAs have been observed in both hands of PD patients at Hoehn and Yahr (HY) stage-I, even when clinical signs are restricted to one side of the body (Hoehn and Yahr, 1967). Additionally, multi-muscle synergies that stabilise posture have been shown to be impaired in patients at HY stage-II, before clinical signs emerge.

The sensitivity of stability control indices to treatment interventions is also well-documented. For example, both synergy index and ASAs show improvement in PD patients undergoing dopamine replacement therapy, such as with L-DOPA (Park et al., 2014; Falaki et al., 2017). Notably, the effects of the first dose of L-DOPA differ significantly from those seen in patients on chronic L-DOPA treatment, who are tested in "off-drug" and "on-drug" states (De Freitas et al., 2020a,b). These findings suggest that chronic exposure to L-DOPA alters the neural circuitry's sensitivity to drug administration and withdrawal. Deep brain stimulation (DBS) is another treatment in PD, and its effects on stability and agility have been mixed. For instance, DBS applied to the subthalamic nucleus or globus pallidus has been shown to improve ASA indices without significantly affecting stability, which could lead to increased falls despite improved agility (Falaki et al., 2018). These observations raise concerns about the widespread use of DBS, especially in cases where agility improves but stability does not, potentially leading to further instability (Marconi et al., 2008; Rocchi et al., 2012; Cossu and Pau, 2017; Bjerknes et al., 2021).

Bernstein's work (Bernstein, 1947a) emphasized the role of subcortical loops, particularly those involving the basal ganglia and thalamus, in the control of synergies. This has been corroborated by more recent studies exploring the involvement of cortico-basal-thalamo-cortical and cortico-cerebellar-thalamo-cortical loops in motor coordination (Thach et al., 1992; Houk et al., 1996). The neurophysiological mechanisms underlying synergies are likely distributed across the central nervous system, with significant contributions from spinal circuits. For instance, segmen-

tal reflexes and the stretch reflex loop are thought to stabilise peripheral movement components (Latash, 2008a). A recent study examined the synergic control of the flexor digitorum superficialis during single-finger force production tasks and found robust force-stabilising synergies, though individual motor unit modes were not stabilised (Madarshahian and Latash, 2021; Aeles et al., 2020). This aligns with work showing a trade-off of features between synergies at different hierarchical levels (Gorniak et al., 2007, 2009).

Lastly, it is important to clarify that the research presented in this thesis focuses exclusively on the development of motor synergy in healthy individuals. In this context, synergy refers to the acquisition of coordinated motor patterns that are optimised through practice to benefit the performance of motor sequence learning. The term "synergy" has a distinct history in the clinical domain, as noted above, where it often describes compensatory motor patterns that emerge following neurological insult, such as stroke (Dewald et al., 1995; McPherson and Dewald, 2022). Clinical synergies, while adaptive in response to injury or dysfunction, frequently lead to compensatory movements as the motor system adjusts to the loss of healthy function. This thesis discusses clinical synergy generally in Chapter 6, but never specifically with experimental data.

1.10 Thesis Overview

This thesis investigates the development of motor synergies in motor sequence learning, focusing on when synergies emerge and evolve during different learning phases. It explores the role of online and offline consolidation in synergy formation, aiming to determine whether synergies are primarily developed over active practice or rest. Additionally, the thesis examines whether synergies formed over early-stage learning of simpler tasks flexibly transfer to more complex tasks.

1.10.1 Methods for the Remote Collection of Motor Data

In Chapter 2, I detail the methods and procedures developed specifically for remote collection of behavioural and motor sequence learning data used in Studies 1 and 2. These studies required adaptations due to the limitations imposed by the COVID-19

pandemic, leading to the creation of a novel data collection framework that allowed participants to perform sequential motor tasks from the safety of their homes. Using a combination of high-definition, high-frame-rate cameras and online experimental software, participants were able to practice a 5-element motor sequence (typing the numbers 4-1-3-2-4 on a standard keyboard) over multiple sessions, while their performance and learning progress were tracked remotely.

1.10.2 Study 1: A comparison of motor synergy development in early and late stages in motor sequence learning

In Chapter 3, I investigated the development of motor synergy over learning a 5-element motor sequence. The study was conducted remotely, with 20 healthy participants over two training sessions. The first session occurred on Day 1, and the second session followed 24 hours later. The primary aim of the study was to examine the rate at which synergy behaviours developed across the early and the late stages of learning. Synergy was quantified by examining changes in motor behaviour across these trials, particularly as Jensen-Shannon divergence (JSD) measures between behavioural distributions. I hypothesized that the bulk of synergy development would occur in the early stage of learning, evidenced by significant changes in JSD during early stage trials, while minimal changes would occur after early stage trials. Moreover, I expected skill gains, measured as the number of correct sequences per second, to correlate with the changes in JSD during early learning but not during the later stage, suggesting a plateau in both performance and synergy development.

1.10.3 Study 2: A comparison of motor synergy development in practice and rest periods in motor sequence learning

In Chapter 4, I present data from 20 healthy participants (also collected remotely) who learned the same 5-element sequence (4-1-3-2-4) over 36 trials of practice interleaved with periods of rest. The primary aim was to investigate the influence of rest periods on the development of motor synergies, specifically whether offline rest or online practice periods accounted for the majority of synergy development. The

methods for data collection were identical to Study 1. I hypothesized that synergy development would occur primarily over rest periods within early stage trials, while significantly fewer changes would occur over online practice periods. An alternative hypothesis was that synergies would primarily form during active practice, with rest contributing to the stabilisation of these patterns rather than further development.

1.10.4 Study 3: A comparison of motor synergy transfer across skills in three durations of training

In Chapter 5, I investigate whether synergies formed during a simple motor skill transfer to more complex skills. Sixty healthy participants were divided into six conditions and trained on a 3-element motor sequence (e.g., typing 4-1-3), followed by a paired 5-element sequence (e.g., typing 2-4-1-3-4), wherein the initial 3-element sequence was embedded. Participants practised the unique 3-element sequence for either 3, 7, or 11 trials and then practised the paired sequence for an equal number of trials. The primary aim was to determine whether synergies developed during the simpler sequence would transfer to the more complex task. I hypothesized that transfer of synergy behaviour would occur after 3 trials, but not after 7 or 11 trials.

1.11 Impact of the Coronavirus Pandemic

The COVID-19 pandemic significantly altered the design of my project. Notably, Chapters 3 and 4 for this thesis were originally designed for a laboratory setting and in-person data collection. However, as the pandemic progressed and restrictions were put in place, in-lab research became infeasible. This led to a re-envisioning of my methods, culminating with a redesign for remote data collection.

With a dedicated support team consisting of Margaret Hayward, Rawan Fakhreddine, and Giselle Rodriguez, from the National Institutes of Health (NIH), National Institute of Neurological Disorders and Stroke (NINDS), and the Human Cortical Physiology Section, we recruited 90 participants who completed the protocol from the safety of their homes. The research was facilitated through the use of high-definition, high frame rate cameras and online experimental software designed to capture fine motor learning details in a remote setting. Data collection involved

capturing video and performance metrics from participants performing motor sequence skills in a manner that closely replicated the intended in-lab studies.

Despite the successful pivot to remote experimentation, the process of setting up remote data collection was not without its challenges. Developing and then awaiting a decision on the novel remote data collection ethics policy took significant time, and the separation from co-workers slowed the preparation required for remote study execution. As noted, Studies 1 and 2, corresponding to Chapters 3 and 4 of this thesis, were most affected by the pandemic. In response to these delays, Chapter 2 of this thesis is dedicated to a comprehensive description of the methods, technologies, both physical and digital, and the data collection pipelines developed for Chapters 3 and 4.

Chapter 2

METHODS AND THE USE OF HIERARCHICAL UNSUPERVISED CLUSTERING OF MARKERLESS POSE DATA FOR DISCOVERY OF MULTI-EFFECTOR SYNERGY

"Measure twice, cut once."

— John Florio

This chapter provides a complete guide to the remote data collection methods, both technologies and procedures, of Studies 1 and 2. Since development of these methods was extensive, and given that they were the foundation for two studies, an account of my methodological reasoning is warranted.

Any methodological variations or additions specific to a particular study will be addressed within the corresponding experimental chapter. For example, Study 3 involved additional equipment and protocols that deviated from the remote framework used in Studies 1 and 2 and as such will be detailed in Chapter 5. Further, the data analysis pipeline was largely consistent across all studies, but any deviations specific to any study will be expanded upon in the relevant chapter's methods.

While the primary aim of this chapter is to detail the preparation and procedural standardization that enabled reliable remote data collection, the secondary aim of this chapter is to detail alternative methodologies available in behavioural annotation, pose estimation, and physical technologies, offering a comparative look between the tools and techniques employed for these research studies and those currently shaping the study of behavioural quantification.

2.1 Advances in Behavioural Quantification

Unsupervised behavioural categorization is a relatively recent development in neuroscience, yet there is a rapidly growing array of tools and techniques available for use. DeepLabCut (Mathis et al., 2018) has become a popular method for defining the positions of body parts in each frame of a video, facilitating easy limb tracking over time. While static limb coordinates are not behaviours in themselves, several tools have been developed to translate these positions into meaningful behavioural categories. Some of these include Deepbehaviour for automated analysis of behavioural kinematics (Arac et al., 2019), VAME for segmenting behaviour using latent space representations and hidden Markov models (Luxem et al., 2020), MoSeq for mapping sub-second structure in mouse behaviour (Wiltschko et al., 2015), B-SOID for unsupervised behaviour identification (Hsu and Yttri, 2021), and Motion Mapper for mapping stereotyped behaviour in fruit flies (Berman et al., 2014). Each

approach has its own advantages and limitations, which need to be carefully considered depending on the specific research question and application.

Table 2.1: Comparison of popular unsupervised behavioural analysis tools

Method	Animal/Human Focus	Temporal Scale	Strengths	Limitations
DeepLabCut	Both (Tracking only)	Frame-by-frame	High-accuracy body part tracking; flexible	Not behaviour classification
DeepBehaviour	Human-focused	Sub-second to multi-second	Kinematic feature extraction; clinical alignment	Requires supervised labels; limited in unsupervised modes
VAME	Animal-focused	Sub-second	Learns latent structure; HMM segmentation	Sensitive to latent dimensionality and tuning
MoSeq	Mice	100–500 ms (sub-second)	Robust to noise; reveals microstructure of behaviour	Highly species-specific; camera and setup dependent
B-SOID	Animal-focused	Millisecond-to-second	Fast, interpretable clusters; minimal preprocessing	Limited validation in human data
Motion Mapper	Fruit fly, mouse	Sub-second to second	Frequency decomposition; no feature engineering	Challenging to scale to high-DoF human motion
HUB-DT	Animal, extended to human	Trial-level or continuous	Unsupervised discovery; adaptable to different species	Originally animal-focused; requires parameter tuning

In this context, behaviour is often treated as a broad and nonspecific concept, with what qualifies as behaviour varying depending on the approach used. This thesis advocates for a machine learning-driven method capable of exhaustively characterising human movements across closely related motor skills. By refining how we define and analyse motor behaviour, we can move beyond subjective classifications toward more objective, data-driven characterizations that allow for a deeper understanding of human behaviour.

To that end, I used HUB-DT (Hierarchical Unsupervised Behavioural Discov-

ery Toolbox) as the primary tool for performing unsupervised behavioural classification (Lindsay and Seamans, 2024). Originally designed for analysing free-ranging animal behaviour, HUB-DT's flexibility made it an ideal choice for capturing and analysing human hand movements during motor sequence learning. Its capacity for unsupervised discovery enabled the identification of motor synergies at a trial-by-trial resolution, offering valuable insights into the timescale and evolution of motor synergy development throughout motor sequence learning.

Despite the rapid advancement of these tools, their adoption in human motor learning research remains limited. This is due in part to several challenges: the higher dimensionality and variability of human movements compared to constrained laboratory animal behaviour, the lack of annotated datasets for supervised training, and the difficulty of validating unsupervised outputs against gold-standard clinical or behavioural metrics. Additionally, some tools - particularly those developed for animal posture - require retraining or substantial parameter tuning to apply meaningfully to fine-grained human motor skills.

HUB-DT supports both traditional behavioural annotation and parameter-free discovery, making it ideal for capturing subtle variations in hand movements. Similar to Motion Mapper, it analyses frequency components of movement through wavelet fitting and applies clustering techniques to classify unique movement patterns. This enabled us to label distinct behaviours during skill acquisition and compare them with manually annotated data, providing a comprehensive understanding of motor synergy emergence and development.

2.2 Behavioural Annotation Software

In the field of behavioural annotation, there are numerous tools, each designed for specific use cases. As noted, HUB-DT was chosen as the primary tool for the purposes of this thesis due to its flexibility in capturing fine-grained motor behaviours during motor sequence learning, as well as its capacity for unsupervised behavioural classification. However, it is important to acknowledge that HUB-DT represents one approach within a diverse landscape of tools. Below is a detailed examination

of several prominent alternatives, each with distinct methodologies, strengths, and limitations. By understanding the range of available tools, I highlight the rationale behind my choice to use HUB-DT and explore how other methods could potentially serve researchers with different data or experimental designs.

Table 2.2: Comparison of behavioural annotation tools by relevance to thesis

Tool	Unsupervised	Variable Behaviour	Low Compute	Fine-Grained Motion	Human-Adaptable
DEEPbehaviour	Optional	✓	✗	✓	✓
Motion Mapper	✓	✗ (periodic focus)	✓	✗	✗
VAME	✓	✓	✗	✓	✗
B-SOID	✓	✗ (short bouts)	✓	✓ (frame wise)	Partial
MoSeq	✓	✓	✗	✓ (3D)	✗
HUB-DT	✓	✓	✓	✓	✓

Ease of use is a major consideration in the adoption of behavioural annotation tools, especially in human motor learning, where experimental pipelines often span modalities. Many tools were developed for high-frame-rate, low-dimensional animal behaviour and require technical adaptation for human data. Additionally, several tools demand high-end GPU compute and extensive tuning, limiting their accessibility for labs without deep ML expertise. These practical barriers contribute to their underutilization in human research, despite their analytical potential.

DEEPbehaviour (Arac et al., 2019) is an early machine learning-based pipeline for end-to-end kinematic analysis. Using a multi-camera setup, it tracks body parts without markers by leveraging convolutional neural networks and long-short-term memory (LSTM) recurrent neural networks. These networks generate bounding boxes for limbs or digits, transforming the data into trajectories. A key aspect of DEEPbehaviour is its flexibility in clustering behaviour, with options for both supervised and unsupervised learning. The dynamic time alignment kernel technique allows for flexible clustering of behavioural trajectories. However, the tool's complexity lies in its computational requirements. Significant time and resources are required for neural network training, and the clustering and classification algorithms need fine-tuning for each experiment. This makes DEEPbehaviour powerful yet highly demanding in customized experimental setups.

Motion Mapper (Berman et al., 2014) focuses on identifying stereotyped behaviours using spectrographic projection. The tool was originally designed for fruit flies and uses multiple camera views to extract posture data, which is transformed into high-dimensional behavioural space using frequency analysis and Morlet wavelets. Dimensionality of the space is then reduced using t-distributed Stochastic Neighbour Embedding (t-SNE). The final step uses the watershed transform to cluster behaviours. Motion Mapper excels at capturing periodic behaviours, which are stereotypical and often frequency-locked. However, this reliance on periodicity can be problematic for non-stereotyped or highly variable behaviours, such as those observed in rodents. Additionally, the method's non-parametric clustering is sensitive to poorly separated density spaces, which can limit its effectiveness in less constrained environments.

VAME (Luxem et al., 2020) uses variational autoencoders (VAEs) to produce a latent behavioural space. The architecture is built on recurrent neural networks that take fixed-length sequences of positional data as input, whether from video annotations or marker-based tracking. The autoencoders compress this data into a lower-dimensional representation, generating a behavioural latent space. Hidden Markov Models (HMMs) are then used to classify behaviours, treating the behavioural space as a dynamical system with discrete states. VAME is particularly useful for detecting short, repeated motifs in locomotion or grooming, making it ideal for open-field experiments. However, the tool requires considerable computational power due to the complexity of training both neural networks and HMMs. Its reliance on fixed-length sequences can introduce bias towards short behaviours, and interpreting the results can be challenging due to the purely probabilistic nature of the HMMs.

B-SOID (Hsu and Yttri, 2021) is an unsupervised tool designed for rodents in open-field experiments. It uses frame-difference measures to track derivatives of position, angle, and relative pose, which are then reduced to two dimensions using Uniform Manifold Approximation and Projection (UMAP) for clustering. The core of B-SOID's clustering method is HDBSCAN, a hierarchical density-based clustering algorithm that identifies clusters in the reduced space as potential behaviours.

B-SOID is well-suited for experiments with repeated behaviours, as it is optimized for high framerates and repeated short-sequence data. However, this specialization limits its effectiveness with smaller, more diverse datasets where local behavioural structures are critical. Adjustments may also be necessary to optimise performance at lower framerates, as noise or slow video capture can reduce accuracy.

MoSeq (Wiltschko et al., 2015) is unique among the tools discussed because it uses 3-D depth imaging to track behaviour without the need for multiple cameras. MoSeq segments behaviour into sub-second motifs, or behavioural "syllables," using autoregressive Hidden Markov Models (AR-HMMs). Like Motion Mapper, MoSeq uses wavelet transformations to quantize behaviours, but it adds the AR-HMM to focus on transitions between behaviours. This probabilistic approach is effective for identifying short, transient actions, providing an additional layer of analysis on behavioural transitions. However, working with 3-D depth data requires careful setup of the experimental environment, particularly if electrophysiology is also being recorded. Additionally, MoSeq's computational complexity and parameter dependency may require substantial adaptation for non-standard use cases, and its focus on transitions means that the definition of behaviours is somewhat dependent on the AR-HMM model rather than the structure of the behaviour itself.

Each of these tools has distinct advantages and challenges, and their effectiveness depends on the specific research question, data characteristics, technological constraints, and experimental design. While DEEPbehaviour and VAME offer flexibility with machine learning-based models, they demand significant computational resources and fine-tuning, which may not be feasible in all settings. Motion Mapper and MoSeq provide robust methods for capturing stereotypical, periodic behaviours but may struggle with more variable or complex actions. B-SOID stands out for its simplicity and effectiveness in experiments with repeated behaviours but may require modifications for diverse datasets.

Ultimately, the choice of behavioural annotation software must be tailored to the requirements of the research. Factors such as the behaviours studied, the available computational resources, and the specific objectives of the experiment are con-

siderations for the most appropriate tool. In this case, HUB-DT was selected for its ability to capture fine-grained motor behaviours during motor sequence learning and its capacity for unsupervised behavioural classification (Lindsay and Seamans, 2024). The next critical consideration in behavioural analysis is selecting appropriate pose estimation software. In the next section, I describe the rationale for my choice of pose estimation software and explore alternative tools.

2.3 Pose Estimation Software

Central to this thesis was the need to select pose estimation software capable of capturing detailed human movement across diverse contexts. After evaluating various computer vision tools, DeepLabCut was the optimal choice due to its accuracy in tracking complex motor behaviours. This section introduces DeepLabCut and then reviews several prominent alternatives, highlighting their methodologies, strengths, and potential applications. Specifically how DeepLabcut was employed in my studies is covered in section 2.5.1, Markerless pose estimation with DeepLabCut.

DeepLabCut (Mathis et al., 2018) is a widely-used tool for markerless pose estimation that applies deep learning to track animal movements. Using a convolutional neural network (CNN), DeepLabCut identifies body parts in videos without the need for physical markers. It was originally designed for lab animals like mice but has since been expanded to other species and even human pose tracking. Its flexibility, combined with the ease of use, makes it a popular choice in behavioural neuroscience and biomechanics. However, like many deep learning tools, its accuracy depends on careful manual labeling during the training phase, and it may require significant computational resources for large-scale datasets.

OpenPose (Cao et al., 2021) is known for its ability to track multiple individuals in real time. This open-source framework uses a bottom-up approach, first detecting body parts and then associating them with individuals. OpenPose can estimate poses for entire bodies, including hands and faces, in both 2-D and 3-D. Its strength lies in its real-time performance, which makes it ideal for applications in sports analysis, video games, and virtual reality. However, OpenPose is relatively

resource-intensive, especially when tracking multiple subjects in high-resolution video, and may require powerful GPUs for optimal performance.

AlphaPose (Fang et al., 2022) is designed to handle real-time multi-person pose tracking. AlphaPose stands out for its high accuracy and ability to address occlusion challenges, where body parts are obscured in video frames. By using a more refined part detection network and an advanced pose recovery system, AlphaPose achieves high performance in detecting body poses, even in crowded scenes. While extremely accurate, AlphaPose is computationally demanding, particularly in real-time applications, but it has been successfully applied in diverse fields like action recognition, behavioural tracking, and human-computer interaction.

PoseNet (Kendall et al., 2015) is a deep learning-based pose estimation algorithm that estimates camera poses in real-time. It was originally designed for mobile device applications, using a lightweight architecture that enables pose estimation without requiring large computational resources. PoseNet performs well for tasks such as augmented reality (AR) and robotics, where understanding the spatial relationship between a camera and the world is critical. Despite its efficiency, PoseNet is limited to estimating 2-D pose and lacks the advanced features found in more sophisticated pose estimation systems.

HRNet (Cheng et al., 2020) is designed to maintain high-resolution representations of inputs throughout its pipeline. Unlike previous methods that downsample input images for efficiency, HRNet keeps high-resolution inputs to preserve spatial details. This results in significantly higher accuracy in pose estimation, particularly in challenging scenarios like complex backgrounds and crowded scenes. HRNet's architecture is highly flexible and can be extended to tasks like object detection and semantic segmentation. Its main drawback is the computational demand, particularly during training, as it requires large amounts of memory and processing power.

ArtTrack (Insafutdinov et al., 2017) is designed for multi-person pose estimation in complex environments. Unlike other methods, ArtTrack focuses on both pose estimation and tracking across frames, making it particularly useful for maintaining identity over time. ArtTrack combines both CNNs and graphical models to

enhance tracking accuracy. However, its complex architecture can make it challenging to implement in real-time without significant computational resources.

DensePose (Güler et al., 2018) extends keypoint detection with dense surface mapping for human body pose estimation. Using a fully convolutional network, DensePose maps every pixel of a detected person to a 3-D surface model of the body, providing a highly detailed and continuous understanding of human pose. DensePose is particularly useful in computer graphics, virtual reality, and human motion analysis. DensePose is computationally demanding, however, and requires sophisticated data collection and processing tools.

Table 2.3: Comparison of pose estimation tools by key functional criteria

Feature	DeepLabCut	OpenPose	AlphaPose	PoseNet	HRNet	ArtTrack	DensePose
Species	Animal, Human	Human	Human	Human	Human	Human	Human
Real-time	✗	✓	✓	✓	✗	● (limited)	✗
Multi-subject	✗	✓	✓	✗	✗	✓	✓
3D pose	● (3D possible)	✓	✓	✗	✓ (high-res 2D)	✓	✓ (dense 3D)
Compute demand	Moderate	High	High	Low	Very High	High	Very High

Each of these tools offers distinct advantages depending on the research aims and experimental requirements. Tools like DeepLabCut and AlphaPose are better suited for highly accurate, animal-specific pose tracking, while tools like OpenPose and DensePose enable real-time, multi-subject tracking in human-centric applications. The trade-offs between computational demand, ease of use, and accuracy make choosing the right pose estimation tool a matter of balancing these factors with the needs of the research.

2.4 Methods for the Remote Collection of Pose Data

Having selected the technologies best suited for my project, I next adapted these tools for remote data collection - a necessity brought about by the COVID-19 pandemic. Traditionally, motor learning experiments are conducted in person using specialized equipment, such as digital tablets, robotic devices, and VR displays, to capture high-resolution temporal and spatial data. However, the pandemic dis-

rupted any capacity for in-person research, introducing new logistical challenges. In this section, I detail my method for the capture of remote, online behavioural data, collected from participant homes amidst lockdown orders.

Remote, online experiments have long been used in the social sciences to gather behavioural data. Platforms like Amazon Mechanical Turk enable researchers to recruit large, diverse participant pools, run pilot studies quickly, and collect data using a range of experimental designs. Compared to in-lab studies, these online methods offer a more representative sample of the general population and can even include individuals with mobility impairments who may otherwise be unable or unwilling to visit a research lab. While the use of online platforms is well-established in many areas of research, there has been limited application in sensorimotor learning studies.

Despite this limitation, existing research suggests that data collected through online studies are often comparable to those obtained in-person. However, prior work in motor learning at home has focused on more naturalistic or ecological setups, limiting the generalisability of findings. To address this, a hybrid platform was developed for the studies in this thesis, integrating the advantages of online research with controlled, lab-based elements. This approach enabled real-time observation and the use of standardised input devices to ensure consistent data collection. In this section, I exhaustively detail my approach, or manual of operating procedure (MOP), for Studies 1 and 2, which is then followed by the data processing pipeline, which, among other components, details my uses of DeepLabCut and HUB-DT.

Pre-Screening

To begin, all potential participants were pre-screened over the phone according to protocol. The responsible investigator entered the responses to the screening questions in the online regulatory binder. Participants were asked questions regarding the inclusion criteria for Studies 1 and 2: if they spoke English, were right-handed, had internet access and a USB port, and if they could type with their left hand without pain. Additionally, eligible participants were required to be between 18 and 90 years old, live in the United States, and not be staff members of the National

Institutes of Health. If these questions were answered in the positive, the potential participant was deemed eligible and enrolled in the protocol.

Enrollment and Study Preparation

Eligible participants received an email with a link to an "Acknowledgement of Participation" agreement (an NIH-approved, low-risk, remote equivalent of a consent form). Once completed, the participant was assigned a unique subject ID code. The participant was then asked if they were able and willing to participate in a second day of the study for additional compensation.

Upon receiving the participant's name and address, home address labels were generated for shipping and return for study materials. The investigator prepared and shipped the materials, updating the tracking spreadsheet accordingly. A confirmation email was sent to the participant with the shipping details and expected delivery date. Upon receipt, a video conference was scheduled with the participant to review the study protocol, confirm the experiment's details, and review all study materials.

Materials Description

Materials used for the remote collection of data included one AKASO EK7000 camera, one Cherry-brand 104 key full-size keyboard, two SanDisk SD cards, one extra camera battery, a battery charger, and a return shipping label. The choice of the AKASO EK7000 camera was motivated by its price-to-quality ratio, offering high-resolution recording capabilities at an affordable price point. The camera was capable of 120 frames per second (fps) at a 720p resolution recordings, which was deemed more than sufficient for capturing the fine-grained movement data necessary for pose estimation. Additionally, the extra battery and charger were requisite, as the study spanned two days, and battery depletion overnight was a concern. The Cherry-brand keyboard was selected due to its precision and reliability, requiring approximately 0.45 N of force for keypress actuation, an average amount that would be consistent with most participants' typing experiences.

On receipt of study materials, participants visually inspected the AKASO EK7000 for any damage. The battery was located, charged. Labels for the cam-

era, battery, and charger were verified by the investigator to ensure identification. After charging, the battery was inserted into the camera, and functionality was confirmed by powering the camera on and off. The SanDisk SD card was also formatted and its serial number was verified. A test photo was taken to ensure the SD card saved data as expected. The keyboard was carefully inspected, plugged into the participant's home computer to verify its functionality, and its serial number was also verified. Lastly, camera settings, lighting, and framing of the hand and keyboard were adjusted to ensure a preferred resolution and clarity. Once all items were properly prepared and double-checked for functionality, the investigator instructed the participant in the set-up of all materials.

Experimental Sessions

Once setup was complete, the participant received the link to the online motor sequence task, carefully reviewed instructions with the investigator, began the camera recording, and then began the motor sequence task. The investigator remained available for technical support but did not interfere with task performance. After the session, the investigator ensured the data was recorded correctly, and the equipment was stored safely for future use or returned as needed.

For participants who agreed to continue to Day 2, the equipment was left in place overnight, and the study was repeated the following day using the same setup. If not, the equipment was dismantled and safely stored.

Post-Participation Activities

After the final session, the participant was responsible for repacking materials and attaching the return shipping label. The investigator arranged for FedEx pickup or drop-off and ensured all study materials were returned promptly. If multiple participants were involved at the same location, the investigator repeated the setup and screening process for each individual before collecting and returning the equipment.

Material Receipt and Inspection

Upon receiving returned materials, the investigator verified that all equipment labels were intact. The AKASO 5700 camera was unpacked, inspected for damage, and

sanitized. The camera's functionality was checked by powering it on, reviewing the participant's videos, and powering it off. The SD card was removed, data was transferred to the NIH server, and the card was reformatted and reinserted into the camera. The camera battery was sanitized, charged, and replaced before the camera was repackaged for further use or storage.

The keyboard underwent a similar process: it was unpacked, sanitized, tested for functionality, and repackaged for the next participant. All materials were then placed back into the parcel box, which included stickers, return labels, and other necessary supplies. If the equipment was not immediately shipped to the next participant, the investigator made note of when the preparation steps were completed. If more than seven days passed before shipment, the camera battery charge was checked again before mailing.

Remote Study Protocol Overview

1: Pre-Screening. Participants were screened via phone to confirm eligibility (age, handedness, language, internet access, and exclusion criteria such as NIH staff status).

2: Enrollment and Consent. Eligible participants signed an NIH-approved "Acknowledgement of Participation" and assigned a subject ID.

3: Shipping. Study kits containing all required materials (camera, keyboard, SD cards, battery, charger, return label) were packed and shipped. Tracking information was recorded and sent to the participant.

4: Setup. On receipt, experimenters guided participants through inspection, charging, testing, and preparing equipment. The camera was positioned and settings were adjusted to ensure hand and keyboard visibility.

5: Experimental Session(s). Participants completed the motor sequence task while recording video locally. The investigator remained on-call but did not interfere. If the participant consented to Day 2, the same procedure was repeated the following day.

To ensure the integrity of the recorded video data, all files were reviewed for completeness and quality prior to analysis. Sessions were inspected for corrupted files, missing footage, or frame rate inconsistencies. Frame counts were verified to confirm 120 fps sampling, and videos were checked to ensure full capture of

the keyboard and participant's hand movements. Trials in which recording failed, the hand was obstructed, or typing deviated substantially from the instructed pattern were flagged and excluded from downstream analysis. These quality control steps ensured that pose estimation and behavioural labelling were based on clean, interpretable input data.

2.5 Online Task and Performance Data

Following equipment setup, participants were emailed a link for the motor sequence task wherein they repetitively typed a 5-item numerical sequence (4-1-3-2-4) as quickly and as accurately as possible. Keypresses were performed with the participant's non-dominant, left hand and applied to a standard, 104-key keyboard. Keypress 1 was performed with the little finger, keypress 2 with the ring finger, keypress 3 with the middle finger and keypress 4 with the index finger (See Figure 2.1). Individual keypress times and identities were recorded for behavioural data analysis. Participants practised the task for thirty-six (36), 10 s duration trials. 10 s rest periods were interleaved between trials. During practice, participants were instructed to fixate on the five-item sequence displayed for the full duration of the trial. Small asterisks appeared above a sequence item when a keypress was recorded, providing feedback to the participant about their current location within the sequence. After completion of a full iteration of the sequence, the asterisks were removed. The asterisk display did not provide error feedback since they appeared for both correct and incorrect keypresses. During the 10 s interleaved rest periods, the sequence was replaced with the letters "REST", which participants were instructed to fixate on. Visual stimuli and task instructions were presented, and keypress responses were recorded using custom code hosted online at Psytoolkit. Psytoolkit is an online platform designed for running cognitive-psychological experiments and surveys. It provides researchers with tools for both programming and administering a wide variety of tasks in behavioural research. For this study, Psytoolkit was used to design and execute the motor sequence task, ensuring consistency in how instructions and feedback were presented. Psytoolkit also recorded a range of performance met-

rics, including both accurate and inaccurate keypresses, as well as time stamps to the millisecond. From these metrics, the primary measure of learning - correct sequences per second (cs/s) - was calculated, providing a quantitative assessment of participants' skill learning over time.

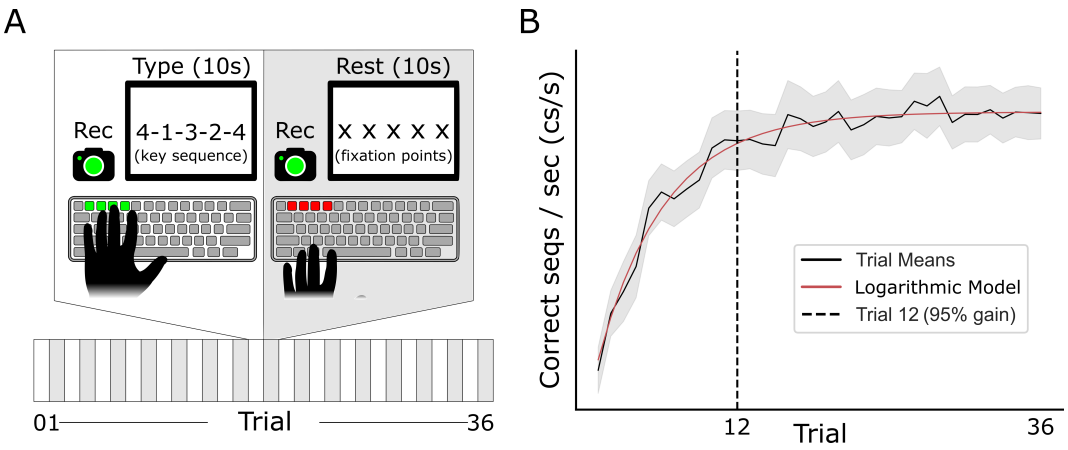

Figure 2.1: Motor Sequence Task and Skill Calculation. (A) Subjects learned a motor sequence over 36 trials within a single training session (Walker, Brakefield et al. 2002, Censor, Horovitz et al. 2014). Each trial consisted of alternating rest and training periods for 10 s duration each. (B) Skill was measured as the typing speed of correct sequences per second (cs/s). Early learning developed over the initial trials of training (trial set within which 95% of total learning occurred, mean ± STD (Bonstrup, Iturrate et al. 2019).

2.5.1 Skill Measure Calculation

Skill learning in this study was defined by a composite measure of speed and accuracy, specifically quantified using the metric of correct sequences per second (cs/s). This metric reflects the participant's ability to accurately execute the entire 5-element sequence within a given time, providing a clear measure of performance at the sequence level. The rate of skill acquisition was modelled using a logarithmic growth curve applied to the cs/s data. By fitting a logarithmic model to these data, I was able to estimate the rate of improvement over time, capturing the gains in performance that typically characterize early motor learning (See Figure 2.1 Right).

To ensure precise identification of correct sequences, individual keypresses were analysed in relation to the specific 5-item sequence "41324." Keypresses were labelled as part of a correct sequence if they matched any circular permutation of this sequence (e.g., "41324," "32441," etc.). This method allowed flexibility in

identifying correct sequences while maintaining stringent criteria for sequence accuracy. By accounting for circular shifts, participants were credited for executing the correct sequence regardless of where they began typing within it, ensuring accurate measurement of skill execution. For each trial, performance was summarized as the mean instantaneous correct sequence speed, providing a detailed measure of how efficiently and accurately participants could produce the required sequences throughout the practice trials.

2.5.2 Learning Stage Calculation

Early-stage learning was defined as the period spanning from Trial 01 up to a cutoff trial, T, at which 95% of the group-level performance gains were achieved. The cutoff trial T was determined through the following process. First, the trial-by-trial group average of correct sequences per second (cs/s) was modeled using an exponential growth function of the form:

$$L(t) = C_1 + C_2 \cdot \left(1 - e^{-\lambda t}\right)$$

where $L(t)$ represents the group average learning state at practice trial t, C_1 accounts for pre-training performance, and C_2 represents the asymptote, or the maximum extent of learning. The parameter λ controls the learning rate, dictating how quickly participants approach their maximum performance.

The parameters C_1, C_2, and λ were estimated using a constrained nonlinear least-squares method, implemented via MATLAB's `lsqcurvefit` function, which employed the trust-region-reflective algorithm. The boundary constraints for these parameters were set as follows: C_1 within the range [0, 5], C_2 within [0, 15], and λ within [0, 2] (Buch et al., 2021).

Next, I accumulated the values of $L(t)$ across all trials t, and divided each trial's cumulative sum by the total area under the curve. This allowed us to compute the trial-by-trial cumulative percentage of learning. The cutoff trial T was then defined as the first trial where the cumulative percentage of learning reached 95%.

2.5.3 Skill Calculation Following Overnight Consolidation

In Study 1, I accounted for the experience of a "warm-up" period of performance, wherein motor sequence skill is quickly regained and eventually matches the level achieved by the end of the previous day. To select the first trial following warm-up on Day 2, I used the same model at which 95% of performance gains are achieved at the group level on Day 1. The model allowed us to pinpoint the first trial following the rapid gains in performance following a warm-up at the group level.

2.5.4 Skill Calculation Following Online and Offline Periods

In Study 2, I assessed early learning changes at the individual trial level by quantifying performance improvements occurring during individual practice and rest periods. Micro-online learning was defined as the change in correct sequences per second between the first and final seconds of each trial. Micro-offline learning was defined as the absolute difference in cs/s between the last second of a trial and the first second of the following trial. These epoch boundaries were chosen to match the resolution of the skill metric (cs/s), ensuring sensitivity to 1 s temporal resolution of skill measurement. Cumulative micro-online and micro-offline learning scores measured over the first T trials for each participant were used to assess their respective contribution to total early learning(Bonstrup et al., 2019).

2.6 Analysis Pipeline for Pose Data

In the following section, I outline the technologies and parameters selected for Studies 1 and 2 with a focus on two key tools: DeepLabCut (DLC) for pose estimation and HUB-DT for behavioural annotation. I will provide a detailed description of how each was applied, along with the specific modifications and extensions made to their underlying codebases to tailor them to the unique demands of my project. Please see Figure 2.2 for a visual of the full pipeline.

2.6.1 Markerless Pose Estimation with DeepLabCut

DeepLabCut (DLC), a machine-learning toolbox designed for pose estimation, was utilized to track the movement of four fingers during the motor sequence keypress

task. DLC employs deep neural networks and transfer learning to precisely determine the position of specific body parts in each video frame. This software has been extensively used for animal tracking but can be applied to human movement analysis. The primary advantage of DLC is its ability to learn the spatial relationships between body parts from labelled training data and generalize these relationships to new, unseen frames during prediction (Mathis et al., 2018).

For Studies 1 and 2, sixty video frames from each participant were randomly selected, resulting in a total of 2,400 frames, balanced across the 36 trials for each study. Frames were manually labelled with the position of the four digits, specifically above the lunula. Labelled frames were then used to train a neural network that generalized digit poses across all frames. The neural network employed in this study was a ResNet-50 architecture, pre-trained on the ImageNet dataset. The pre-training allowed the network to leverage a large body of prior knowledge about image features, improving the accuracy of pose estimation. Training was performed using the DLC toolbox (available at https://github.com/DeepLabCut/DeepLabCut).

Training involved one million iterations and utilized three colour channels (RGB), with pairwise terms, but without additional supervision. Data augmentation was performed before training, as necessary, by re-scaling images to maintain a consistent image scale across all frames. This augmentation process helped increase the network's robustness to variations in image conditions, ensuring reliable tracking. After training, the neural network was used to predict the position of the four fingers in all frames of every trial. The predicted position of a finger in each frame was determined by the peak of the output score-map generated by the network. Frames where fingers were occluded or not clearly visible were identified by low peak scores, indicating uncertainty in the network's predictions.

Training and subsequent predictions were conducted on the NIH High Performance Computing Biowulf cluster, allowing for efficient processing of the large dataset. Validation of the network's tracking accuracy was conducted post hoc. This was achieved through visual inspection of predicted finger trajectories for each participant, along with the use of a likelihood parameter provided by DLC. The

likelihood parameter, which quantifies the confidence of the network's predictions, was manually set at 95%. Pose estimation below this threshold was inspected manually and either included or excluded from the dataset, based on its accuracy. Missing data points were linearly interpolated to maintain continuity in the tracking data. Pilot tests indicated that 60 manually labelled training frames were sufficient to train a network that could estimate finger positions across all frames and participants with an accuracy greater than 90%.

Output from DLC analyses consists of CSV files. In Studies 1 and 2, each CSV contained approximately 43,200 rows (86,400 for all data over two days). Rows correspond to the x, y positional data of each digit, per frame, sampled at 120 Hz over a duration of 720 seconds (the total experiment time) minus 360 seconds (the total rest period time). These data were used for subsequent kinematic analysis to investigate the emergence of motor synergy during learning of the motor sequence.

2.6.2 Morlet Wavelet Transformation of Pose Data

DLC output provides frame-by-frame 2D coordinates for each feature of interest. While these instantaneous body part locations are valuable, they do not inherently represent behaviour. Behaviour, as we define it, arises from changes in pose that unfold across varying timescales. Thus, to capture behaviour effectively, both spatial (position) and temporal components must be considered, particularly at multiple temporal scales. Traditional methods that rely on fixed time periods - common in many existing behavioural analysis tools - present several key limitations. These are discussed extensively in the Motion Mapper paper by Berman et al. Notably, fixed-time period methods constrain temporal information to a single scale, which often leads to the omission of multi-scale interactions across and within features - interactions that are fundamental to many typical behaviours.

The raw DLC traces do contain some temporal information, such as rates of change in positions. However, these traces are limited in their ability to distinguish between behaviours with similar positional changes but differing dynamics. For example, a rapid oscillatory movement and a slow, smooth displacement might produce similar x, y displacements over a given period but represent entirely different

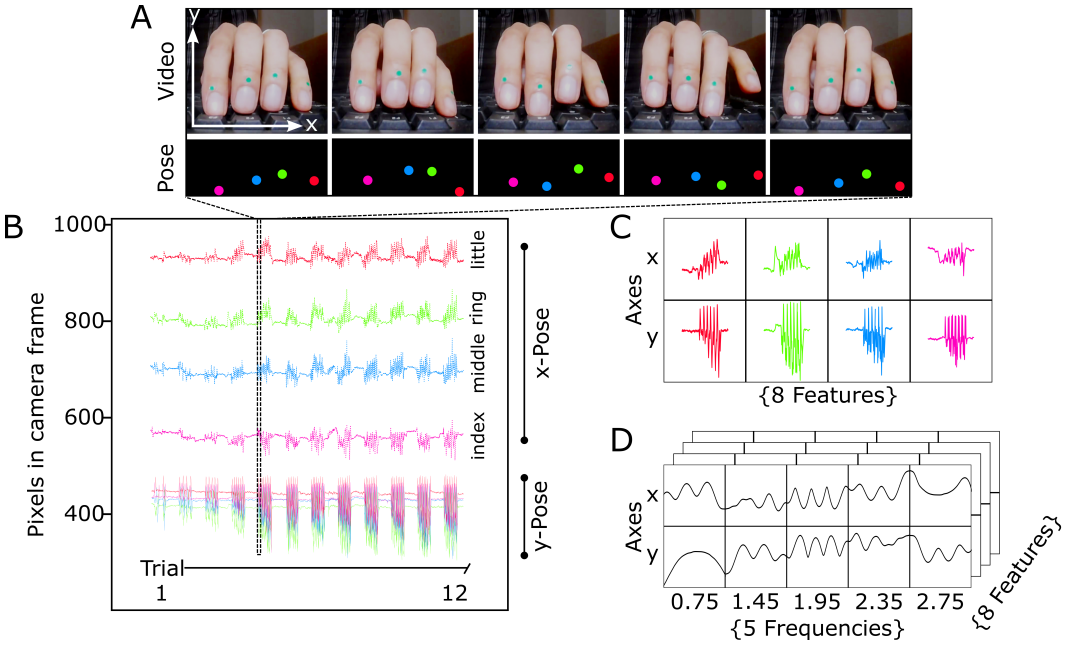

Figure 2.2: Pose Data and Morlet Wavelet Transformation. (A) Top: High-definition video was used for markerless tracking of digit movements. Bottom: Markerless pose estimation identified x, y poses of each digit. (B) Poses identified from markerless tracking depict digit positions during performance of keypresses in the motor sequence task. Finger displacements were calculated along the x and y axis of the camera frame in terms of pixels. Trials 1-12 are shown. Tracked features are transformed from the (C) temporal domain to the (D) frequency domain by applying the Fast Fourier Transform. The transformed features are then convolved with a series of Morlet wavelets, at widths set by the frequencies over which I aim to capture behaviour. The amplitude of response to the wavelets is then transformed back to the temporal domain by the inverse Fast Fourier Transform, resulting in a set of wavelet amplitudes. Performing wavelet projection on the whole set of features produces a full set of wavelet responses, which form the behavioural space whose dimensionality is expanded by a fold equal to the number of chosen frequencies.

behaviours. To address this, I used HUB-DT's spectrographic projection with Morlet wavelets to transform the raw spatial coordinates into an expanded space that simultaneously captures both spatial and temporal features at multiple timescales. Morlet wavelets, which are complex Gaussian functions modulated by a sinusoidal wave, are particularly effective for highlighting periodic and dynamic patterns in the data, such as the oscillations seen in chunking of keypress sequences. This transformation enables the representation of motion patterns across varying temporal frequencies, effectively disambiguating behaviours that might otherwise appear similar in the raw traces.

The resulting dataset embeds each frame of video into a high-dimensional spa-

tiotemporal manifold. The dimensions of this manifold correspond to the response amplitudes of each feature (and its two spatial coordinates) to each wavelet frequency used in the projection. For my analysis, the manifold spans 5 wavelet frequencies, 4 features of interest, and 2 tracked coordinates per feature, providing a rich and detailed description of the motion of tracked features.

While this manifold offers a comprehensive description of the tracked features, its high dimensionality introduces practical challenges. High-dimensional spaces often include redundant or irrelevant information, which can impair the effectiveness of clustering algorithms - a phenomenon known as the "curse of dimensionality" (Bellman, 1957, 1961; Hastie et al., 2011). To address this, I performed dimensionality reduction before clustering. This step helps distil the manifold into a more computationally manageable representation while preserving the most informative aspects of the data. Although algorithms like neural networks can learn latent representations directly from high-dimensional data, I prioritised interpretability by explicitly reducing dimensionality before unsupervised clustering.

2.6.3 Fit Wavelet Projections into a 2-D Embedding

Dimensionality reduction is a key step in data pre-processing and is often initiated using Principal Component Analysis (PCA). PCA creates a sequence of uncorrelated projections, which are linear approximations of the input data, ordered by variance. A subset of these projections can be selected to form a lower-dimensional representation of the data, effectively preserving most of the information. This reduced space is particularly useful for identifying global patterns and trends in the data. However, for behavioural data, where key information may be embedded in local patterns, PCA may not always be the optimal approach. Several other methods of dimensionality reduction are more suited to preserving local structures. At the start of my project, I had two primary options.

t-SNE (t-distributed Stochastic Neighbour Embedding) is an alternative to PCA that emphasizes the preservation of local relationships in the data. t-SNE produces low-dimensional embeddings (often in 2 or 3 dimensions) that maintain proximity relationships from the original high-dimensional space. It does this by

modelling similarities between points in both the input and output spaces as probability distributions - Gaussian in the input and a t-distribution in the output. This ensures that local structures, such as clusters and manifold relationships, are retained, though at the expense of global accuracy. The primary parameter in t-SNE is perplexity, which controls the number of nearest neighbours used to preserve local distances. While highly effective, t-SNE is computationally intensive and transductive, meaning it cannot project new data into an existing embedding. Recent advancements, such as those from van der Maaten (2014) and Linderman et al. (2017), have improved its efficiency through methods like neighbourhood approximation and interpolation, as implemented in the openTSNE package.

UMAP (Uniform Manifold Approximation and Projection) offers another powerful approach to dimensionality reduction, which shares similarities with t-SNE but follows a distinct methodology. UMAP begins by constructing a fuzzy topological representation - modelling data relationships probabilistically to account for uncertainties and preserve both local and global structures - of the data that captures local relationships in its high-dimensional form (McInnes et al., 2018). This topological structure is then optimized to fit a lower-dimensional space using cross-entropy, creating an embedding that reflects the original data's structure. UMAP is parametrised by the number of neighbours, which adjusts the scale of local relationships, and it includes a minimum distance parameter that controls the spacing of points in the low-dimensional space. These parameters allow for fine-tuning and greater flexibility in preserving local structures. UMAP is computationally efficient and scalable, making it suitable for large datasets and applications requiring rapid dimensionality reduction.

Ultimately, I chose HUB-DT's UMAP implementation over t-SNE for this work due to several critical advantages, particularly in terms of computational efficiency and scalability. While t-SNE offers excellent local structure preservation, it is computationally expensive, especially for large datasets. Given the repetitive nature of human hand movements during sequential typing, where I needed to track subtle changes in behaviour over time, UMAP's faster runtime and abil-

ity to handle larger datasets beneficial. UMAP's ability to preserve both global and local structures, combined with its scalability, allowed us to explore fine-grained behavioural synergy without being limited by the computational overhead associated with t-SNE. Further, UMAP's flexible parameterisation gave us finer control over the reduced space, which was essential for detecting small but significant changes in typing patterns over the course of skill learning.

2.6.4 Hierarchical Unsupervised Clustering

The low-dimensional embeddings of the behavioural space offer a valuable visualisation of behaviour patterns across the entire dataset. These embeddings facilitate the identification of recurrent motion patterns that may exhibit consistent characteristics across trials, i.e. motor synergy. Visualising this structure helps distinguish between stereotyped, repetitive movements and more variable, exploratory actions. The unsupervised segmentation of behaviours, therefore, emerges naturally from clustering techniques applied to this reduced space, allowing us to classify and group behaviours based on their intrinsic characteristics.

To introduce how I clustered behaviour in the reduced space, I'll first describe Density-Based Spatial Clustering of Applications with Noise, or DBSCAN, a well used algorithm for clustering high-dimensional data. DBSCAN operates without making assumptions about the shape or distribution of clusters. Instead, it identifies core samples - points that reside in regions of high density - and expands these regions by including neighbouring points within a specified distance, denoted by the parameter ε (Ester et al., 1996). Clusters are thus formed by expanding from core points, and isolated points that do not meet the density threshold are labelled as noise. The key advantage of DBSCAN lies in its ability to detect clusters of varying shapes and sizes, making it particularly effective for behavioural data that may not conform to predefined cluster shapes.

However, DBSCAN employs a single density threshold (ε) to form clusters, which may oversimplify complex, fine-grained data such as ours, where behaviours occur at multiple levels of granularity. Recognizing this limitation, I opted for an extension of DBSCAN that addresses these issues - Hierarchical Density-Based

Spatial Clustering of Applications with Noise (HDBSCAN).

HDBSCAN extends the principles of DBSCAN by introducing a hierarchical approach that captures clusters across varying densities, making it more suitable for complex behavioural datasets (Campello et al., 2013; McInnes et al., 2017). Rather than relying on a fixed density threshold, HDBSCAN constructs a minimum spanning tree based on reachability distances between points in the data, generating a hierarchy of clusters that persist across different levels of density. This hierarchical structure allows the algorithm to identify stable clusters that exist across multiple scales, eliminating clusters that are too small or too diffuse to be meaningful by applying a minimum cluster size parameter.

In the context of my behavioural data, where movements vary across digits and participants in terms of frequency, duration, and consistency, HDBSCAN's hierarchical nature enables the detection of clusters that may not be captured by simpler algorithms like DBSCAN. By condensing the cluster hierarchy and discarding unstable or noise-laden clusters, HDBSCAN provides a set of robust, well-defined clusters that represent meaningful patterns in the data. This capability is particularly valuable in detecting synergies in motor sequence learning, as it allows us to observe how clusters evolve across trials. HDBSCAN ensures that the identified clusters accurately reflect the underlying structure of the behaviour, offering insight into the progression of motor sequence learning and underlying motor synergies.

Each frame of video is represented as a row of wavelet-transformed pose features, which are projected into a two-dimensional space using UMAP. These 2D coordinates serve as the spatial representation of each frame. HDBSCAN then operates on this low-dimensional space, grouping nearby points into clusters based on local density. Frames that fall into high-density regions are assigned cluster labels, while those in sparse areas are labelled as noise (–1). In this way, each frame receives a label based solely on its location in the UMAP space, allowing for unsupervised segmentation of behaviour across the time series.

Labelled clusters can be visualised by colour coded plots in the embedded space. Each colour indicates a separate cluster, corresponding to a putative be-

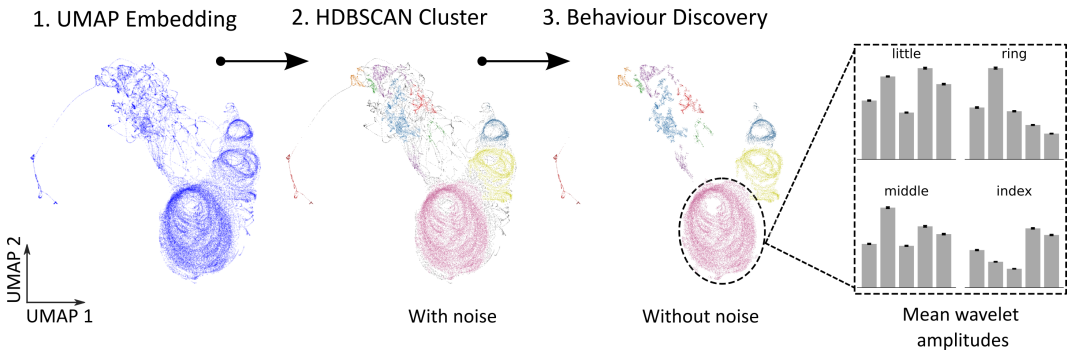

Figure 2.3: Left: A 2D UMAP embedding of high-dimensional pose-derived wavelet data, where each point represents a single video frame. Center: HDBSCAN clustering assigns labels based on local density; noise is shown in low opacity. Right-center: After removing noise, distinct behavioural clusters emerge, including the circled pink cluster of interest. Right: Mean wavelet amplitudes for each digit (little, ring, middle, index) are computed for the circled cluster to characterise its kinematic signature.

haviour. Therefore, each cluster in the label set corresponds to a single discovered behaviour. The pose kinematics of that behaviour can be visualised by plotting the mean wavelet amplitudes for points in a cluster. This is the averaged wavelet response for each feature, and together shows the changes in pose across the range of selected frequencies that characterise that behaiour. Here, these clusters represent the first step in identifying motor synergy.

2.6.5 Synergy Analysis Pipeline

To this point, I have been leveraging the functions available in the HUB-DT package. Following the acquisition of cluster labels from HDBSCAN, I diverged from the HUB-DT package to examine motor synergy in detail. While HUB-DT provided an essential foundation for unsupervised behavioural classification and clustering, it became necessary to develop custom code that would allow for a deeper exploration of the temporal and spatial dynamics of motor synergies. This shift in methodology enabled us to analyse motor behaviours at a finer granularity, focusing specifically on how synergies evolve and stabilize over time within a motor sequence. The custom pipeline described facilitated the identification, quantification, and statistical validation of motor synergies across trials.

Frame-by-Frame Label Tracking and Processing

To capture and analyse motor synergy during skill learning, a custom pipeline was developed to process the frame-by-frame labelled data collected during experiments. Each row of the data represented a frame and the associated motor behaviour label for that specific frame. I then created a function to tally continuous occurrences of each behaviour label. Noise, as determined by the HDBSCAN clustering, was represented by the label -1 and was filtered. For each non-negative label, I tracked the number of continuous frames during which that behaviour occurred. The start frame and total count for each occurrence across trials, as well as the mean duration in frames of each behaviour was recorded to isolate the time-series of motor synergies for analysis.

These extracted features were formatted into a structured array, which served as the foundation for my analysis, including time-series comparisons and the calculation of velocity and acceleration metrics. This method ensured a systematic approach to tracking and analysing motor synergy throughout the experiment.

Label Distribution Across Trials

Following the frame-by-frame label tracking, I calculated and plotted the distribution of labels across all trials, both excluding and including noise labels. I calculated the relative frequency of each label within each trial, normalizing the counts to control for varying amounts of noise. This normalization allowed us to observe the proportional distribution and evolution of motor synergies across trials. This step was necessary to assess the prevalence and frequency of labelled motor synergies during motor sequence learning.

I then repeated this process but included noise labels in the analysis to assess their impact on the data. The distributions were visualized using stacked bar plots, where the x-axis represented trial numbers and the y-axis indicated the proportion of labels present (in other terms, I display the proportion as ‘percent use’ per trial).

Label-Specific Kinematics

To analyse for synergy from each label, I calculated several data points:

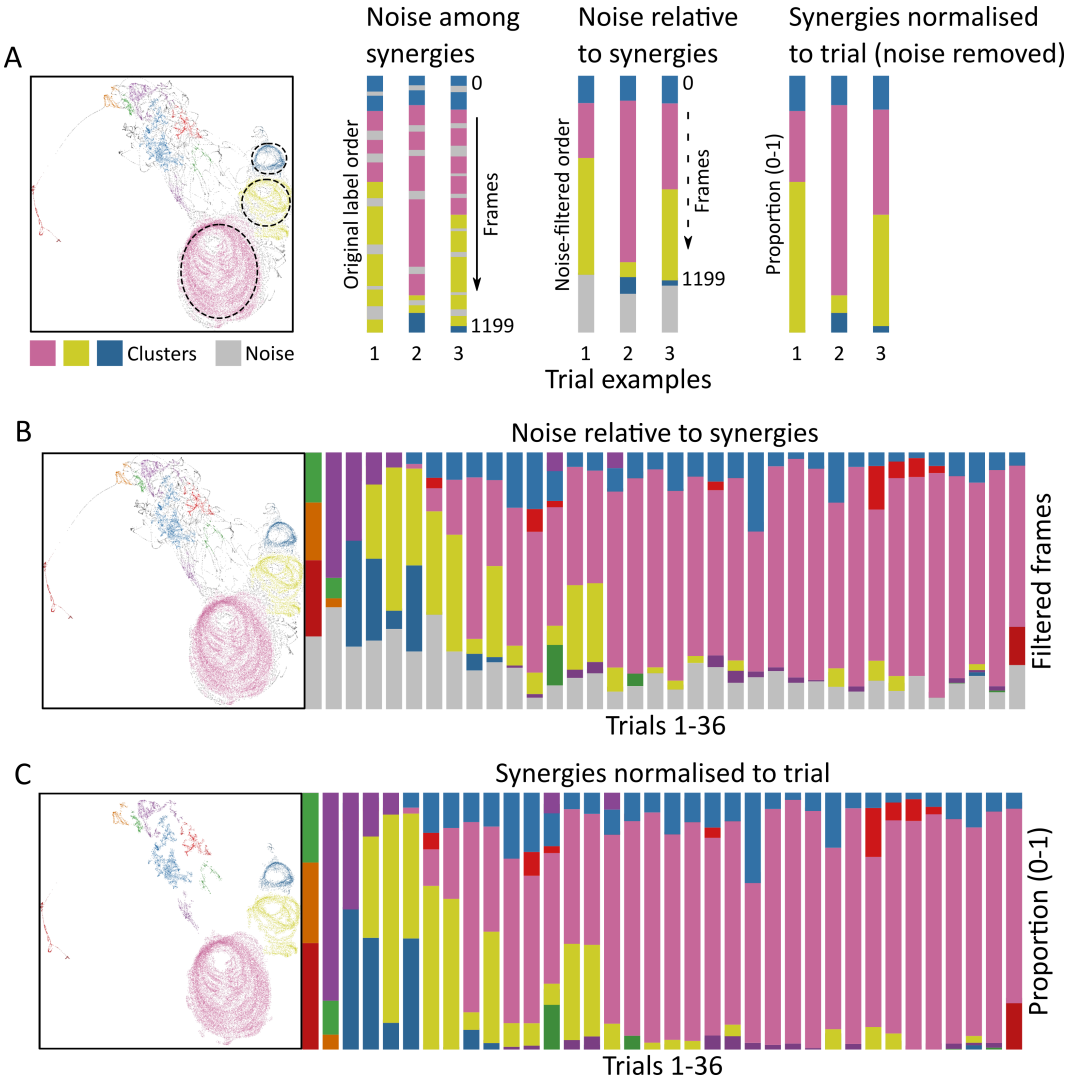

Figure 2.4: (A) Left: A UMAP projection of the wavelet-transformed amplitude data is shown, with HDBSCAN-derived cluster labels visualised as coloured points. Three example synergies are circled for reference. The left-centre portion of the panel displays the original behavioural labels across three real trials, each shown as a stacked bar spanning 1200 video frames. Each frame is assigned a label - colours represent synergy labels, while grey denotes frames labelled as noise. In the right-centre portion, noise labels are filtered and collected at the base of each bar to visualise the proportion of noise within each trial; synergy-labelled frames remain unchanged. Finally, in the rightmost panel, noise is fully removed, and the remaining synergy labels are normalised to now represent all behaviour in a trial as a distribution of synergy behaviours. (B) Shows the behavioural label distributions (including noise) for all 36 trials, with grey indicating noise and colour showing identified synergies. Panel C displays the same distributions after noise has been excluded and synergy behaviours scaled to sum to one per trial. These normalised distributions form the basis for trial-by-trial comparisons of motor synergy development throughout training.

First, I determined the **keypress count**, defined as the number of correct keypresses within each label. This metric provided a direct measure of participant performance within each labelled behaviour. Note in Figure 2.5 B and Figure 2.6 B that keypress events are determined by first calculating the keypress onset and offset events and then detecting the peak of the time series between the onset and offset.

Next, I calculated the mean **duration** of each label in milliseconds. This was achieved by converting the frame count of each label segment into time, using a frame rate of 120 frames per second. By assessing the duration of each label, I could examine the temporal characteristics of motor synergies.

Further, I computed for normalised correct keypresses, referred to as **normalised Hz**. This involved normalising the keypress count within each label to account for variations in label duration. Specifically, I adjusted the keypress frequency to a standard reference duration of 1 second (120 frames). This normalisation was essential to prevent distorted interpretations when comparing labels of different durations. For example, two labels might both have a keypress count of four, but if one label spanned 90 frames and the other 200 frames, their activity levels per unit time would differ significantly. By normalising to Hz, I could accurately compare the activity levels across labels, identifying which labels were more active in terms of correct keypresses within a fixed window. This approach allowed us to evaluate the relative success of different motor synergies in the context of the task.

Calculations of keypress counts, durations, and normalised frequency for each label gave insights into the emergence and evolution of motor synergies as participants practised the keypress task. Results for each label were then compiled and exported to a CSV format. The CSV file also contained 'presence-of-label' or 'absence-of-label' markers (0 for absent, 1 for present) for each trial. The exported CSV data enabled us to evaluate the occurrence and underlying kinematics of motor synergies across participants.

Dynamic Time Warping and Synergy Identification

To identify all occurrences of a synergy from the data, I employed dynamic time warping (DTW), a technique designed to align sequences with temporal shifts. This

algorithm is particularly useful in multi-effector motor synergy identification, as it allows us to compare sequences that may vary in time or speed within a threshold while retaining the same core structure. DTW computes the optimal alignment by minimizing the distance between sequences, enabling the identification of synergy patterns across trials.

I first extracted the target synergy, defined by the start and end frames of its first occurrence. The DTW algorithm was then applied to locate similar patterns of motor behaviour across the remaining trials. This method allowed us to search for patterns that were not perfectly time-synchronized but exhibited a highly similar motor structure. The DTW distance for each comparison was calculated, and the lowest distances were used to select the most similar occurrences of the synergy.

The DTW distance between two time series X and Y, each of length T, is computed by constructing a cost matrix C, where each entry $C(i, j)$ represents the Euclidean distance between points x_i and y_j.

The time series are defined as:

$$X = \{x_1, x_2, \ldots, x_T\}, \quad Y = \{y_1, y_2, \ldots, y_T\}.$$

The goal is to find a warping path W:

$$W = \{w_1, w_2, \ldots, w_K\},$$

where each element

$$w_k = (i_k, j_k)$$

aligns point x_{i_k} with y_{j_k}, such that the total cost is minimized:

$$\mathrm{DTW}(X,Y) = \min_W \sum_{k=1}^{K} C(w_k)$$

This path provides an optimal alignment between the two sequences.

Once each occurrence of the synergy was identified across trials, I plotted the mean and confidence intervals for each digit time series to visualize how reliably

the synergy was performed. The value in plotting these mean and confidence intervals was also to support the claim that the behaviour identified was indeed a multi-effector synergy, rather than isolated or infrequent movements. The aim here is to quantify the emergence and evolution of motor synergies across skill learning.

This interpretive step - from clustered behaviour to candidate synergy - requires methodological justification, particularly in light of traditional dimensionality-reduction approaches. Classical synergy detection methods such as PCA, NNMF, or UCM typically operate on discrete, isolated motion segments (e.g., single reaches) and output as a fixed number of components. By contrast, my approach identifies continuous, multi-effector behaviours that unfold continuously over time, reflecting coordinated patterns that occur within the span of the task - rather than being extracted from isolated movements or predefined event windows.

Rather than assuming coordination through decomposition-first techniques, I begin by extracting wavelet representations of kinematic time series and projecting them into a lower-dimensional space using UMAP. I then apply unsupervised clustering (via HDBSCAN) to this embedding to identify dense, recurring patterns of movement. These candidate synergies are validated through their statistical density, recurrence across trials, and consistency in kinematic execution. This offers a behaviourally grounded and statistically defensible alternative - one that prioritises what the learner actually does, rather than what a model is constrained to extract. In this way, synergy is discovered through emergent regularity and contextual embeddedness, not merely inferred from dimensional compression.

Time Series Analysis and Kinematic Calculations

For each identified synergy, time-series data for each digit were extracted and analysed to compute key behavioural metrics. Kinematic data were processed to determine the onset, peak, and offset of all digit movements leading to keypresses. Specifically, the Euclidean distance $d(x,y)$ between successive data points x_i and x_{i+1} in the time-series was calculated to capture changes in digit position over time:

$$d(x_i, x_{i+1}) = \sqrt{\sum_{j=1}^{n} (x_{i,j} - x_{i+1,j})^2}$$

where n is the number of dimensions (typically x, y, and z coordinates) and j indexes each dimension. From these distances, I identified the onset, peak, and offset of motor events for each digit.

Velocity v for each digit was calculated as the rate of change in position over time, using the equation:

$$v = \frac{\Delta d}{\Delta t}$$

where Δd is the change in Euclidean distance between frames and Δt is the time interval between successive frames. Acceleration a, the rate of change in velocity, was then derived by measuring changes in velocity across consecutive frames:

$$a = \frac{\Delta v}{\Delta t}$$

Additionally, I calculated what is known as an overlap metric, which assigns a value to the temporal coordination of multiple digits. Specifically, the overlap metric quantified the percentage of frames during which multiple digits were active at once, i.e. the duration of time wherein the onset of a subsequent digit movement occurred prior to the offset of the preceding digit movement.

The percentage of overlap was calculated as follows:

$$\text{Percent Overlap} = \frac{\text{Number of Overlapping Frames}}{\text{Total Frames}} \times 100$$

By analysing the overlap of digit movements, I could evaluate how digits "worked together" during motor sequence learning. This metric, combined with velocity and acceleration data, provided a 'fingerprint' of the temporal dynamics of any given synergy, contributing to a deeper understanding of how motor synergies evolve and stabilize with practice.

In Figures 2.5 and 2.6, I present visual summaries of the analysis pipeline

applied to pose data. Figure 2.5 illustrates the results for a motor synergy that emerged during early-stage learning and became increasingly prominent over the course of training. In contrast, Figure 2.6 depicts the results for a motor synergy that gradually diminished and became nearly extinct as training progressed. Each figure consists of four elements (A, B, C, D), which are described in detail below.

Element A of Figures 2.5 and 2.6 highlights two aspects of the identified motor synergy. The left panel shows a single HDBSCAN cluster representing all frames labelled with the synergy of interest. Data are drawn from one participant. The right panel displays the distribution of this synergy across 36 trials, with each vertical bar corresponding to a single trial (as described in Figure 2.4). In this context, the grey segments do not represent noise. Rather, they reflect all other synergy labels, which have been desaturated to visually emphasise the synergy of interest. The height of the coloured segment in each bar indicates the proportion of the synergy in that trial, illustrating how its expression changes over training.

Element B shows the y-axis pose of each digit across the duration of a representative instance of the synergy. The x-axis explicitly lists the duration of the synergy in video frames (120 frames = 1 s). Each line corresponds to the y-pose of one digit, and keypress events are annotated: green circles for onsets, red stars for keypresses, and magenta circles for offsets.

Element C shows the average pose (lines) and 95% CI (shaded regions) for each digit across all occurrences of the synergy. These plots were generated by aligning the time series using dynamic time warping (DTW), which accounts for temporal variability between instances. This visualisation captures both the consistency and variability in digit movement patterns that define the synergy.

Element D quantifies three kinematic features of the synergy. The first panel presents the mean velocity (pixels/second) for each digit; the second shows mean acceleration (pixels/s^2); and the third measures the percentage of temporal overlap between digits, capturing a degree of coordination between digits over the synergy.

It is important to note that although clustering was unsupervised and therefore performed without temporal linearity, HDBSCAN clusters do capture statistically

distinct, recurring patterns of digit coordination that span linear behaviours. As noted, these patterns - some encompassing full 5-element sequences, others representing common sub-sequences or transitional fragments - are treated as candidate motor synergies. Importantly, their definition is based on spatial-temporal structure, not on frame-by-frame adjacency. This approach enables the discovery of behavioural motifs that recur across trials, even if they do not occur at precisely the same moment in each trial. Figures 2.5 and 2.6 illustrate this principle, showing how the frequency of specific synergies rises or falls across the 36 training trials. These dynamics - emergence, stabilisation, or extinction of synergies - underscore their behavioural relevance and functional role in learning.

Given this framework, I used Jensen-Shannon Divergence (JSD) to quantify how the distribution of behavioural strategies shifts across trials. Rather than comparing individual synergies directly, JSD compares the full distribution of behavioural labels within each trial, capturing how learners change their reliance on different motor patterns over time. Further, JSD treats label distributions as unordered sets, meaning that the numerical label identities (e.g., cluster 0, 1) and their order have no bearing on the divergence - as only the relative frequencies matter. This design is sensitive to the re-weighting of strategy use shown in motor skill acquisition, especially when behaviours are phased out or replaced. The choice for JSD is justified in Section 2.6.6, where I explain its robustness to sparsity, boundedness, and symmetry - critical properties for comparing probabilistic data.

2.6.6 Jensen-Shannon Divergence, Permutation Test, and Stouffer's Method Calculations

To quantitatively assess the development of motor synergies across participants, I detail exactly how I extract label distributions, then detail steps to compute Jensen-Shannon Divergence (JSD) values between trials, and statistically evaluate these comparisons through permutation tests and p-value adjustments.

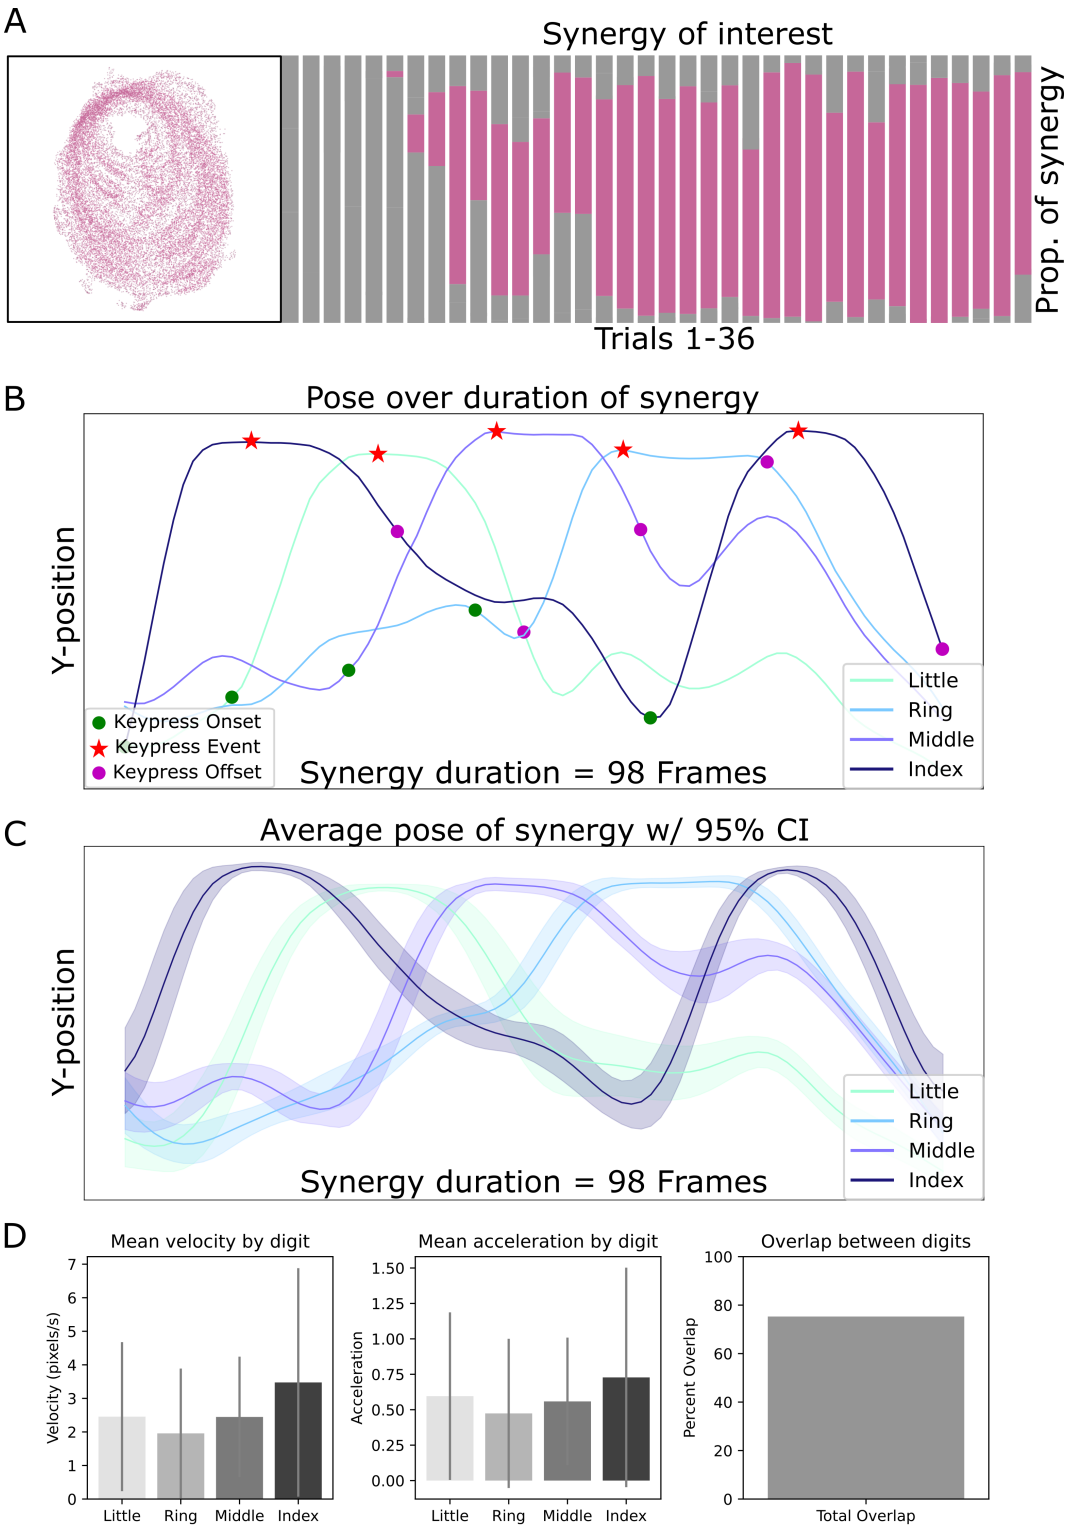

Figure 2.5: Synergy analysis of a representative participant, Example A.

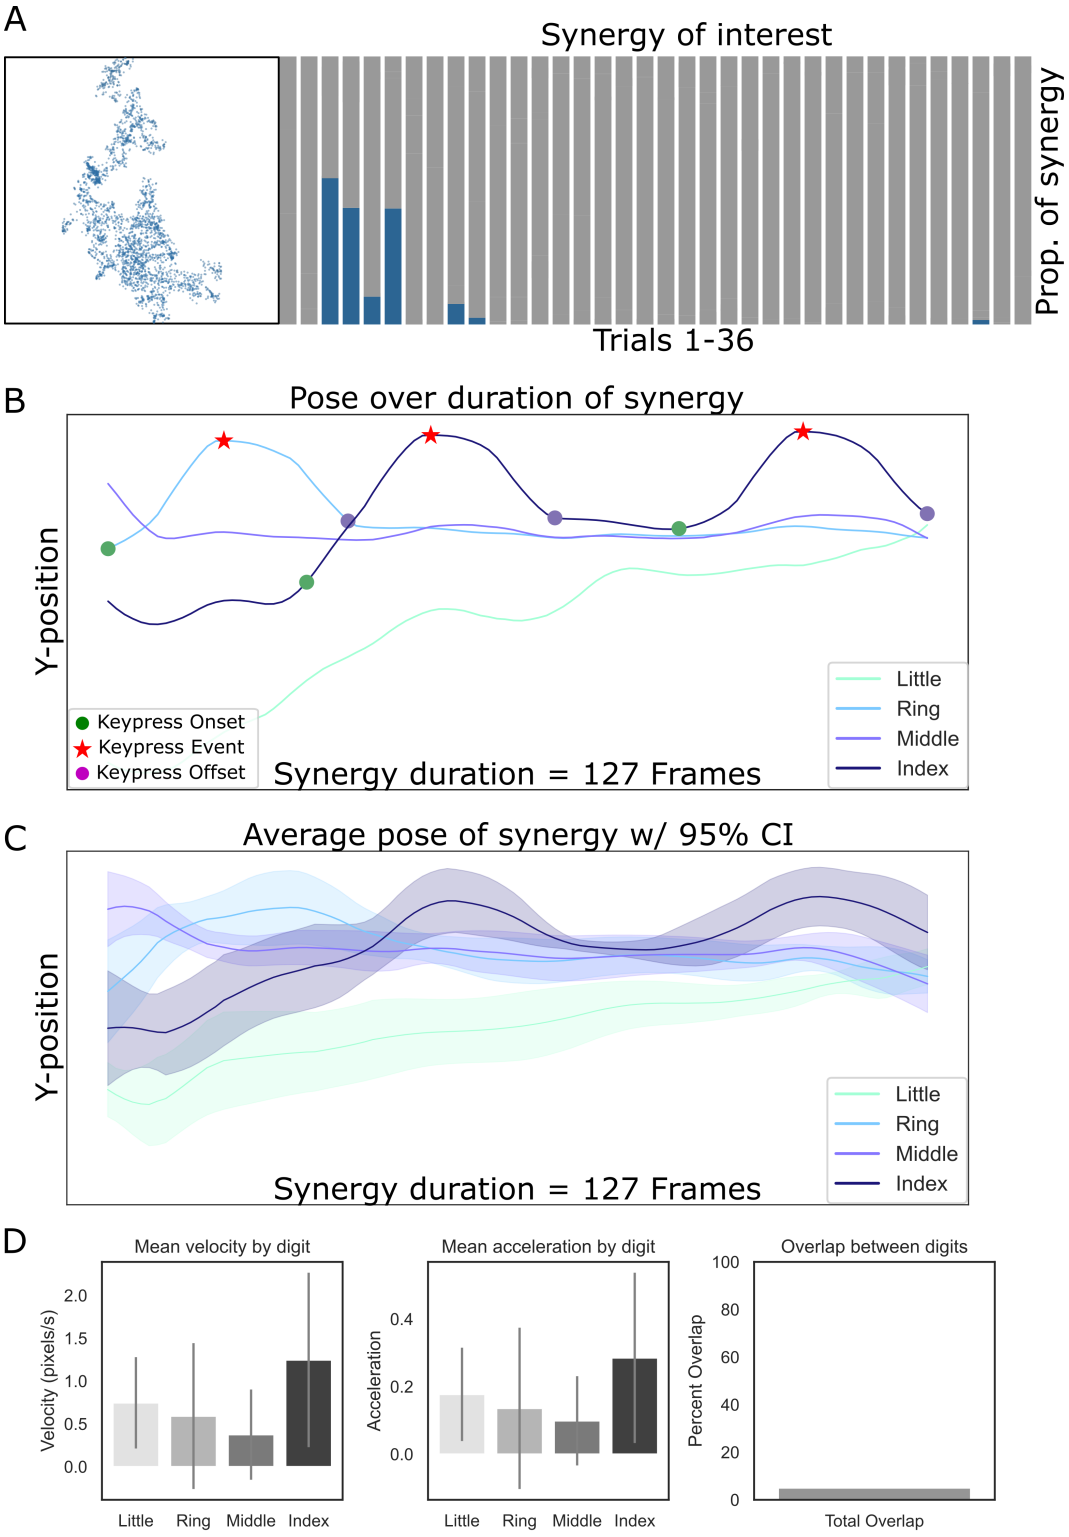

Figure 2.6: Synergy analysis of a representative participant, Example B.

Data Extraction and Preprocessing

Behavioural labels from each participant were stored in `.spydata` files. These files were systematically processed by a custom Python function to ensure consistency in analysis. The first step involved loading the `.spydata` files into a dictionary format for each participant:

```
def load_spydata(file_path):
    try:
        data, error = io.load_dictionary(file_path)
        if error:
            raise ValueError(f"Error loading {file_path}: {
                error}")
        return data
    except Exception as e:
        print(f"Failed to load {file_path}: {e}")
        return None
```

Label Distribution Comparison

For each trial, behavioural clusters were compared by counting the number of occurrences of each label and converting these counts into probability distributions. These distributions were then used to calculate the JSD between different trials.

```
def compare_clusters(cluster_labels1, cluster_labels2,
    all_labels):
    counts1 = np.array([np.sum(cluster_labels1 == label) for
        label in all_labels if label != -1])
    counts2 = np.array([np.sum(cluster_labels2 == label) for
        label in all_labels if label != -1])
    return counts1, counts2
```

The comparison function processed each `.spydata` file, extracted the behavioural labels, and converted these labels into counts and proportions.

JSD Calculation

JSD was calculated between the probability distributions (proportional distributions of synergy) of different trials. JSD measures the similarity between two distributions and is expressed as:

$$\mathrm{JSD}(P\|Q) = \sqrt{\frac{1}{2}\left(\sum_k P(k)\log\frac{2P(k)}{P(k)+Q(k)} + \sum_k Q(k)\log\frac{2Q(k)}{P(k)+Q(k)}\right)}$$

Where P and Q represent the probability distributions of behavioural labels at two different time points, and k represents each cluster label.

Suitability of JSD for Comparing Behavioural Proportions

When comparing proportions of behaviour, selecting an appropriate distance measure is essential. For my purposes, *JSD* has proven to be particularly effective. JSD is a symmetric and smoothed version of the *Kullback-Leibler Divergence (KLD)*, and it offers several advantages over other distance measures. However, to fully appreciate why JSD is the most suitable option for my data, it is crucial to explore both the properties of JSD and the alternative distance measures, evaluating their strengths and limitations.

JSD is favoured for several reasons. First, it is symmetric, ensuring that $JSD(P\|Q) = JSD(Q\|P)$, which is critical when comparing behavioural proportions because it treats the distributions equally without assigning directional bias (Lin, 1991). Second, JSD values are bounded between 0 and 1, with 0 indicating identical distributions and 1 indicating maximal divergence. This bounded range provides a clear and interpretable metric, allowing for consistent comparisons across multiple trials or participants. Third, JSD can handle zero values in the distributions, which is especially important in behavioural data where certain behaviours may be absent in some trials but present in others. This capability avoids the problematic infinities seen in KLD when one distribution has a zero probability for an event (Wong, 2019). Finally, JSD is sensitive to both large and small differences between distributions, which is particularly valuable in motor learning studies where behavioural

changes emerge gradually across trials.

The rationale for using JSD extends beyond its mathematical properties. In this context, JSD was not used to compare individual synergies directly, but rather to compare the full distribution of behavioural labels within each trial. This design reflects my focus on compositional strategy - how learners reallocate effort across a set of motor patterns, rather than substituting one for another. Comparing individual synergies would obscure this broader dynamic, particularly during skill acquisition, where older behaviours are phased out and new ones emerge with different frequencies. JSD is uniquely suited to capturing this re-weighting of motor synergies across time, providing a principled and interpretable measure of change in behavioural strategy.

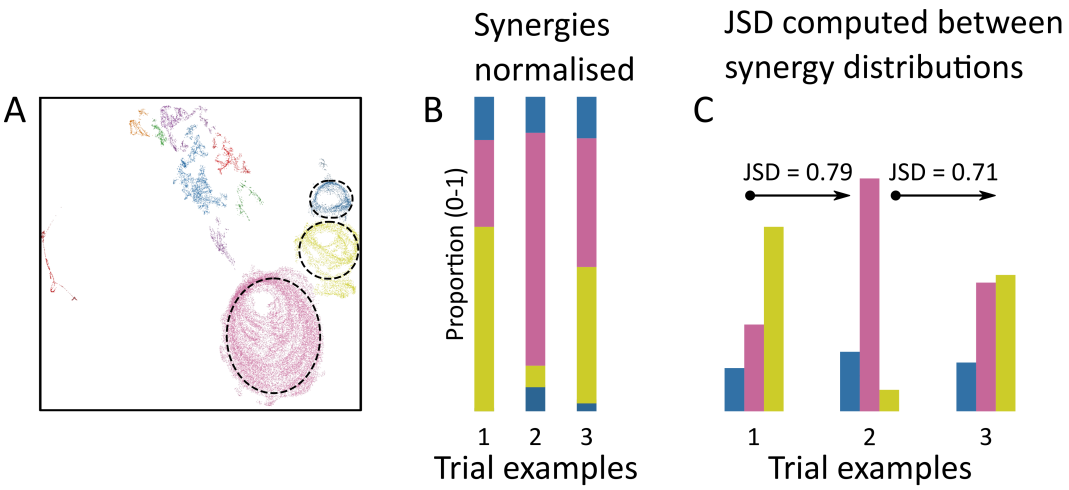

Figure 2.7: Illustration JSD calculation between trials based on normalised synergy labels. (A) UMAP projection of behavioural embedding with selected clusters (synergies) circled. (B) Normalised behavioural label distributions for three example trials (noise removed, proportions scaled to 1). Each vertical bar represents the synergy composition within a single trial. (C) The same data from (B) restructured to highlight the distribution of synergies across trials. Bars of the same colour represent the same label, and their heights reflect the proportion of that label within each trial. When a label occurred multiple times within a trial, its proportions were summed (e.g., the blue label in trials 2, 3). JSD values are calculated between the full label distributions of different trials (e.g., trials 1 vs. 2 and trials 2 vs. 3), providing a quantitative measure of change in behavioural strategy across learning.

Despite JSD's advantages, there are alternative distance measures that could be considered for comparing behavioural proportions, and they may be more appro-

priate for different datasets. The *Euclidean distance*, for example, is a simple and intuitive metric that measures the straight-line distance between two points (Xu and Wunsch, 2003). While it is computationally simple and easy to interpret, Euclidean distance is sensitive to scaling issues and may not appropriately handle probabilistic data like ours, where the magnitude of differences is crucial. This limitation makes it less suitable for behavioural data with categorical distributions.

The *Kullback-Leibler Divergence (KLD)* is another prominent alternative. KLD is widely used in information theory to quantify how one probability distribution diverges from a reference distribution, making it useful in scenarios where the direction of divergence is important (Kullback and Leibler, 1951). However, KLD is asymmetric, meaning $KLD(P\|Q) \neq KLD(Q\|P)$, and this can introduce bias depending on the order of comparison. Additionally, KLD is particularly problematic when dealing with zero values in the distributions, as it tends to diverge to infinity in such cases, making it unsuitable for sparse behavioural data like ours (Wong, 2019).

The *Bhattacharyya distance* measures the overlap between two statistical distributions. It takes into account both the mean and variance of the distributions, providing a more detailed comparison of distributional similarity (Bhattacharyya, 1946). However, it can be computationally intensive, especially for high-dimensional data, and like KLD, it does not handle sparse distributions well.

The *Wasserstein distance* or Earth Mover's Distance measures the cost of transporting probability mass from one distribution to another and is especially effective for comparing distributions that differ in shape (Villani, 2008). This makes Wasserstein distance particularly useful for continuous data, such as spatial distributions, but its computational expense makes it less practical for large datasets like ours. Additionally, while Wasserstein is robust to distributions with non-overlapping support, this property is not necessary for my behavioural data, which tend to exhibit consistent overlap across trials.

Finally, *Cosine similarity* measures the cosine of the angle between two vectors, focusing on the orientation rather than the magnitude of the distributions (Singhal, 2001). Cosine similarity is often used in high-dimensional comparisons, such

as text data, where the direction of data points is more important than their magnitude. However, for my purposes, where the magnitude of differences in behavioural proportions is crucial, cosine similarity is less appropriate.

Therefore, I concluded that JSD is suited for comparing label distributions. Symmetry, boundedness, and ability to handle zero values make it effective for quantifying the overall divergence between synergy distributions across trials, without the need to account for directional differences or continuous distributions.

Permutation Test for JSD Significance

To assess the statistical significance of each observed JSD value, a permutation test was applied. This method shuffled the labels between trials, recalculating the JSD for each permutation to derive an empirical p-value.

```
1  def permutation_test_jsd(proportions1, proportions2,
       observed_jsd, n_permutations=100000):
2      combined_counts = np.concatenate((proportions1,
           proportions2))
3      n = len(proportions1)
4      perm_jsds = []
5
6      for _ in range(n_permutations):
7          np.random.shuffle(combined_counts)
8          perm_proportions1 = combined_counts[:n]
9          perm_proportions2 = combined_counts[n:]
10         perm_jsd = jensenshannon(perm_proportions1,
               perm_proportions2)
11         perm_jsds.append(perm_jsd)
12
13     perm_jsds = np.array(perm_jsds)
14     p_value = np.sum(perm_jsds >= observed_jsd) /
           n_permutations
15     return observed_jsd, p_value
```

The permutation test p-value was calculated as:

$$\text{Permutation P-value} = \frac{\text{Number of permutations with JSD} \geq \text{Observed JSD}}{\text{Total number of permutations}}$$

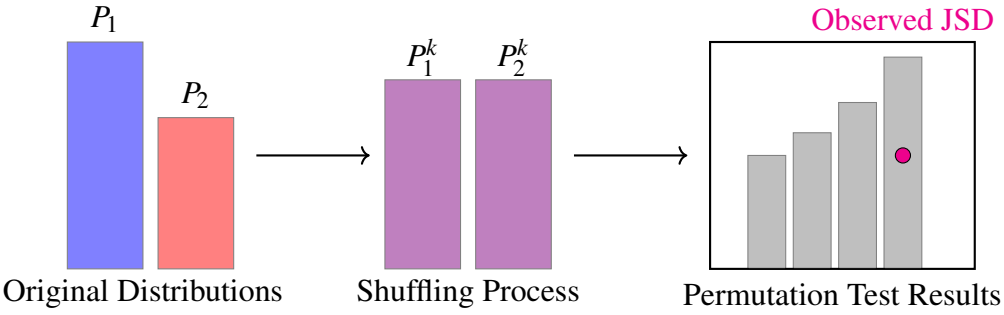

Figure 2.8: Schematic of the permutation test procedure used to evaluate JSD between distributions. (Left): Original distributions P_1 and P_2 represent the behavioural label frequencies for two trials. (Middle): These labels are randomly shuffled and reassigned to generate permuted distributions (P_1^k, P_2^k), simulating the null hypothesis of no difference. (Right): A null distribution of JSD values is built from many such permutations (grey bars), against which the observed JSD (magenta dot) is compared to calculate an empirical p-value.

Combining P-Values Across Participants Using Stouffer's Method

After calculating p-values for each participant, values were combined across the cohort using Stouffer's method. This method aggregates p-values into a combined test statistic. p-values from each participant were first converted into z-scores:

$$z = \Phi^{-1}(1 - p/2)$$

Then, the combined z-score was calculated as:

$$Z = \frac{\sum z}{\sqrt{n}}$$

where n is the number of participants. The combined p-value was determined using:

$$\text{Combined P-value} = 2\,(1 - \Phi(|Z|))$$

This combined p-value reflects the overall significance of the differences in

distributions across participants.

```
def combine_pvalues(p_values):
    z_scores = np.array([norm.ppf(1 - p / 2) for p in
        p_values])
    combined_z = np.sum(z_scores) / np.sqrt(len(z_scores))
    combined_p = 2 * (1 - norm.cdf(np.abs(combined_z)))
    return combined_p
```

The significance of the combined p-value was determined using a predefined significance threshold:

```
def determine_significance(p_value):
    if p_value < 0.001:
        return "p < 0.001"
    elif p_value < 0.01:
        return "p < 0.01"
    elif p_value < 0.05:
        return "p < 0.05"
    else:
        return f"p = {p_value:.3e}"
```

With Stouffer's method, I combined p-values from each participant, ensuring a comprehensive analysis of the overall significance of group differences in motor synergy distributions between trials.

2.7 Methods Summary

In this chapter, I provided a complete overview of the methods used in Studies 1 and 2 of this thesis, focusing on remote data collection and analytical techniques for motor sequence learning. I began by discussing the advances in behavioural quantification, highlighting the rapidly growing array of tools and techniques available for unsupervised behavioural categorization in neuroscience. Emphasizing the need for machine learning-driven methods capable of exhaustively characterising human movements across closely related motor skills, I selected HUB-DT as the

primary tool for unsupervised behavioural classification due to its flexibility and capacity to capture subtle variations in hand movements.

Next, I examined alternative behavioural annotation software, providing a detailed evaluation of several prominent tools such as DEEPbehaviour, Motion Mapper, VAME, B-SOID, and MoSeq. Each tool was detailed in terms of its methodologies, strengths, and limitations, highlighting the rationale behind my choice of HUB-DT and exploring how other methods could potentially serve researchers with different data or experimental designs.

I then addressed the selection of pose estimation software, which was central to analysing human movement across diverse contexts. After evaluating various computer vision tools, I selected DeepLabCut due to its accuracy in tracking complex motor behaviours. I provided an overview of several prominent alternatives, including OpenPose, AlphaPose, PoseNet, HRNet, ArtTrack, and DensePose, discussing their methodologies, strengths, and potential applications.

Following this, I detailed the methods for the remote collection of motor data, a necessity brought about by the COVID-19 pandemic. I adapted traditional in-person motor learning experiments to a remote setting by developing a hybrid platform that integrated the advantages of online research with controlled, lab-based elements. The methods of operating procedure included pre-screening of participants, enrolment and study preparation, and the provision of essential materials such as high-resolution cameras and standardized keyboards. I described the experimental sessions conducted in participants' homes, post-participation activities including the return of equipment, and the procedures for material receipt, ensuring data integrity and consistency throughout the study.

Finally, I described the analysis pipeline for pose data, focusing on the application of DLC for markerless pose estimation and HUB-DT for behavioural annotation. DLC was first employed to extract x and y poses of finger movements from video of the keypress task, generating frame-by-frame coordinates for each digit. These data were transformed using Morlet wavelets to capture both spatial and temporal features of behaviour, projecting the data into a high-dimensional behavioural

space. UMAP was then applied to fit the high dimensional wavelet projections into a two-dimensional embedding, preserving local and global structures. Hierarchical clustering using HDBSCAN identified clusters in the 2-D embedded space. I developed a custom synergy analysis pipeline to analyse clustered behaviours in detail, including calculations of keypress counts, durations, normalized frequencies of keypress counts, as well as velocity, acceleration, and digit overlap. Analyses such as JSD, permutation tests, and Stouffer's method, were conducted to quantitatively assess the development of motor synergies across trials and participants.

In sum, the methods detailed in this chapter were developed to support Studies 1 and 2 of this thesis, enabling reliable remote data collection and analysis of motor sequence learning. The following chapters will present data from Studies 1 and 2, on the timing and development of motor synergy in motor sequence learning.

Table 2.4: Summary of tools and parameters used in Studies 1 and 2

Stage	Details
Pose Estimation	**Tool:** DeepLabCut (DLC) **Parameters:** Tracked 5 digits on left hand; 120 fps, 720p. Custom-trained model per participant using 200 manually labelled frames.
Behavioural Clustering	**Tool:** HUB-DT **Parameters:** Morlet wavelets (25 scales); UMAP (n_neighbours = 30, min_dist = 0.3); HDBSCAN (min_cluster_size = 50).
Dimensionality Reduction	**Tool:** UMAP **Parameters:** Applied to wavelet output. Metric = cosine or Euclidean; 2D output.
Clustering	**Tool:** HDBSCAN **Parameters:** Applied to 2-D UMAP space. min_cluster_size = 50, min_samples = 10. Clusters refined via stability index.
Similarity Metric	**Tool:** Jensen–Shannon Divergence **Use:** Compared synergy distributions across early and late trials.
Statistical Tests	**Tools:** Permutation testing, Stouffer's method **Parameters:** 100,000 shuffles for null distribution; combined p-value via Z-score aggregation.

Chapter 3

MOTOR SYNERGY CHANGES SIGNIFICANTLY WITH PRACTICE AND LARGELY IN EARLY-STAGE LEARNING

"Beginnings are such delicate times."

— Frank Herbert

3.1 Introduction

Motor skill acquisition enables individuals to perform activities ranging from basic daily tasks to complex professional and athletic skills. Early stages of motor sequence learning are characterised by rapid performance improvements as learners transition from isolated movements to coordinated sequences. Understanding the mechanisms underlying these rapid enhancements is crucial for developing effective training and rehabilitation strategies (Shadmehr and Krakauer, 2010).

Sequential motor skills, such as playing a musical instrument or typing, require integrating individual movements into fluid sequences, forming multi-effector motor synergies. These synergies support tasks like grasping, typing, finger spelling,

and haptic exploration (Turvey, 2007).

Although significant performance gains occur during early learning, the specific kinematic mechanisms facilitating these improvements are not fully understood. Previous research suggests that increased movement speed directly benefits skilled performance of motor sequences (Doyon and Ungerleider, 2002), but these studies often focus on simple, discrete tasks involving a single limb and effector, limiting generalisability to real-world scenarios. Understanding how multiple effectors coordinate in continuous motor sequence skills requires analysing the kinematics of each effector over time. The timeline of motor synergy development across effectors during early learning remains unknown, or at the very least, evidence disagrees (Latash, 2008b; Karni et al., 1995).

Early and late stages of motor learning engage different neural substrates. Early learning involves increased activity in brain regions responsible for attention, planning, and error correction, such as the prefrontal cortex, premotor areas, cerebellum, and parietal regions (Hikosaka et al., 2002). As learning progresses, reliance shifts towards the primary motor cortex with decreased activation in prefrontal areas, reflecting skill automatisation and reduced cognitive demands (Poldrack et al., 2005). While my study focuses on behavioural aspects, acknowledging these neurological processes provides context for understanding how stages of motor learning may influence motor synergy development.

Determining the rate at which synergies develop during early learning is essential for tailoring training programmes to maximise motor skill acquisition. If synergy formation occurs rapidly, interventions can enhance early motor transformations through targeted practice strategies. Conversely, if synergies develop gradually, training may need to support ongoing motor synergy emergence and use over extended periods with sustained practice and reinforcement.

Current Study

Study 1 aims to understand the rate at which motor synergies develop during early and late stages of learning a novel motor sequence. I examined changes in motor synergy use across practice trials to quantify synergy development using JSD

between motor synergy distributions. Twenty healthy participants learned a five-element keypress sequence over two training sessions, with the first session on Day 1 and the second 24 hours later.

I hypothesised that multi-effector motor synergy development would be significantly greater during the early stages of learning than during later stages. This expectation is supported by evidence indicating that rapid skill gains occur early in the learning process due to swift neural and behavioural adaptations (Dayan and Cohen, 2011). Additionally, studies have shown that muscle synergies can develop quickly in natural motor behaviours, such as grasping, suggesting that the nervous system efficiently organises motor elements into coordinated patterns during initial practice (d'Avella and Bizzi, 2005; Overduin et al., 2012). While these findings highlight the importance of synergies in manual behaviours requiring coordinated muscle activity, they do not indicate whether such rapid synergy development is possible in continuous motor sequence tasks. I expected minimal changes in synergy formation during later stages, indicating a plateau in both performance and motor pattern optimisation. Furthermore, I anticipated that skill gains - measured as the number of correct sequences performed per second - would correlate with changes in synergy measures during early learning but not during later stages.

By clarifying the timeline of motor synergy development during early learning, this study aims to enhance understanding of the kinematic mechanisms underlying rapid skill acquisition. The findings may inform the design of training and rehabilitation protocols by identifying critical periods when interventions can most effectively promote motor skill development.

Significance and Implications

Understanding the rate and extent of motor synergy development during early learning has significant implications for neuroscience, physical therapy, and skill training. If most synergy formation occurs during initial practice sessions, training programmes can focus on this critical period, incorporating strategies that enhance synergy development. This approach may lead to more efficient learning processes and improved motor performance.

In rehabilitation settings, identifying when motor synergies develop can inform interventions for patients recovering from motor impairments. By targeting early learning phases, therapists may facilitate the re-establishment of functional movement patterns and expedite recovery. Additionally, understanding the plateau in synergy development can help determine when to introduce new challenges or variations in training to promote continued motor adaptation. Overall, this research contributes to our understanding of motor skill acquisition by exploring the temporal dynamics of motor synergy development during early learning.

3.2 Methods

Methods for Study 1 are exhaustively detailed in Chapter 2, but I will provide a brief overview here of those elements most relevant to the proceedings.

3.2.1 Overview

In Study 1, I used high-definition video capture of participant movements through online, remote data collection methods to record performance in a motor sequence learning task. Participants learned to type the sequence 4-1-3-2-4 as quickly and accurately as possible, with each digit specifically assigned to one number: the index digit for number 4, the middle digit for number 3, the ring digit for number 2, and the little digit for number 1. Motor sequence skill learning was captured using the online platform PsyToolkit, with performance quantified as correct sequences per second (cs/s), reflecting both speed and accuracy in sequence execution. The rate of skill acquisition was modelled using a logarithmic growth curve applied to the cs/s data, allowing estimation of improvement over time and capturing the gains characteristic of early stage learning. Digit movements captured in the videos were analysed using DLC to estimate pose data across all trials. Pose data were dimensionally expanded using Morlet wavelet transformations, projecting the data into a time-frequency domain to capture both spatial and temporal features of the movements. Wavelet projections were subsequently fitted into a two-dimensional space using UMAP. With dimensionality reduction and hierarchical clustering via HDBSCAN, I analysed the pose data to discover motor synergy use over trials.

Kinematic analyses were conducted across participants for each digit within each clustered behaviour to determine motor synergies, including calculations of keypress counts, duration of behaviour, normalised frequency of keypresses, velocity, acceleration, and digit overlap. Please take note that use of the phrase 'keypress counts' here is used to illustrate the number of keypresses made within a clustered behaviour and is separate from the accurate or inaccurate keypresses made and counted by Psytoolkit during the online task. Because data were processed individually through UMAP and HDBSCAN - and since there is no practical method to perform embedding alignment on UMAP embeddings - I performed permutation tests on individual participant JSD scores between early and late stage learning trials. To determine whether significant changes in motor synergy use occurred across trials at the group level, I used Stouffer's method to combine p-values from individual analyses. My approach emphasises the necessity for individual data processing and allows for an understanding of motor synergy development across participants.

3.2.2 Inclusion criteria

Participants were recruited if they were between 18 and 90 years old, spoke English, and were right-handed. Additional inclusion criteria required participants to have an internet-connected home computer with a USB port or USB port adapter, be able to participate in a videoconference with a member of the research team, and be capable of typing with their left hand without pain or discomfort. Participants had to be located within the United States to receive study materials. Exclusion criteria included being a staff member of Human Cortical Physiology Section (HCPS; my section with the NIH). If participants were unable to receive materials within the United States, they were required to withdraw from the study or defer participation until they could receive materials in the country.

3.2.3 Participants

Twenty naive right-handed healthy participants (12 women; mean ± SD age 29.2 ± 2.12) gave their written informed consent to participate in the project, which was approved by the Combined Neuroscience Institutional Review Board of the

National Institutes of Health (NIH). The study was split between two days, 24 hours apart from one another. However, at the time of pre-screening and recruitment, 10 of the 20 participants declined to participate on the second day, resulting in only 10 participants completing both sessions. Active musicians were excluded from the study (Abraham and Drory, 2014). The sample size was determined via a power analysis of skill learning data collected using the same task (Censor et al., 2014).

3.2.4 Study design

Study 1 uses a within-subject, repeated-measures experimental design to investigate changes in motor synergy during motor sequence learning over 10 s periods of practice and rest. Participants learned a motor sequence task - typing the sequence 4-1-3-2-4 as quickly and as accurately as possible, with each digit specifically assigned to one number. Over two days, participants completed 36 trials per day, with each trial consisting of 10 s of practice followed by 10 s of rest. The primary outcome measure was the JSD, calculated from pose estimation data to quantify differences in motor synergy distributions between early and late stages of learning. By employing a within-subject design with repeated measurements, the study allowed for direct comparison of each participant's motor performance across different learning stages. This design minimised inter-subject variability and enhanced statistical power, making it well-suited for detecting significant changes in motor synergy use attributable to the learning process.

3.3 Results

In Figure 3.1 (Left), we see skill gains through the early learning stage over Day 1. The black line represents trial means, while the red line shows a fitted logarithmic model that effectively captures the rapid improvement in correct sequence speed (measured in cs/s) during the initial trials, eventually levelling off at 95% saturation by Trial 12 (the dashed, vertical line). This plateau indicates the end of the early learning phase for the motor sequence. The mean correct sequence speed at Trial 1 is 0.472 (SEM = 0.048), increasing to 0.949 (SEM = 0.062) at Trial 12, corresponding to the 95% gain in performance. The gain observed from Trial 1 to Trial 12 is

0.476 ± 0.078, reflecting significant improvements in participants' performance.

Data from Day 2 show a distinct phase of post-warm-up performance. The fitted logarithmic model and trial means indicate that the initial improvement on Day 2 is more modest, as participants quickly return to their prior skill level after an initial warm-up phase. The performance begins at a mean of 0.839 cs/s (SEM = 0.075) in Trial 1 and reaches 1.075 cs/s (SEM = 0.072) by Trial 3, which is marked as the point of 95% skill saturation following the warm-up. The observed gain from Trial 1 to Trial 3 is 0.237 ± 0.104. The reduced variability and plateaued performance beyond Trial 3 indicate a stabilisation of motor skill, suggesting that participants quickly regained proficiency before entering a phase characterised by stable, consistent execution of the learned motor sequence.

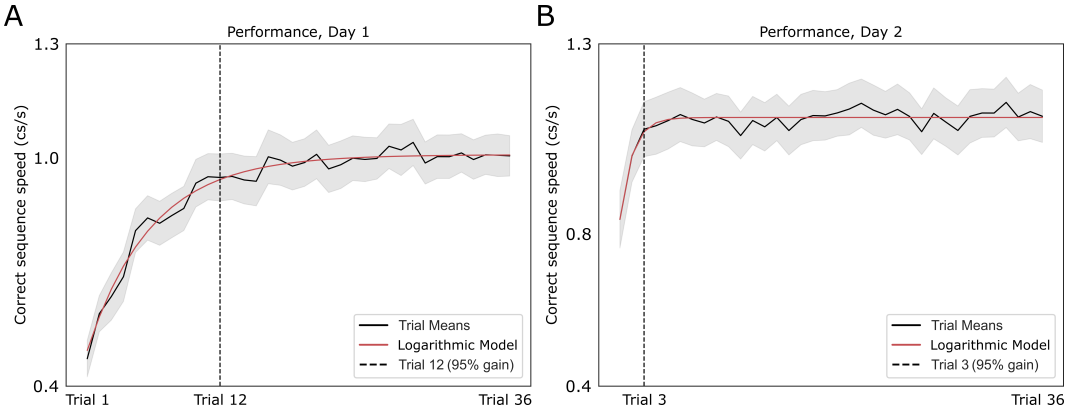

Figure 3.1: Skill was measured in correct sequences per second (cs/s). The black line indicates mean trial performance, and the red line shows the fitted logarithmic model. Vertical dashed lines mark the point of 95% skill gain - Trial 12 on Day 1 and Trial 3 on Day 2. Shaded regions represent the standard error of the mean (SEM) across participants.

Results of the permutation test between Trials 1 and 12 are summarised in Table 3.1, presenting the observed JSD value, corresponding permutation p-values, and Z-scores for each of the 20 participants. Notably, several participants exhibit significant p-values, specifically Participants 001, 003, 006, 018, 019, and 020, with p-values below the conventional threshold of 0.05, suggesting that their observed JSD values are unlikely to have occurred by chance. The corresponding Z-scores for these participants are also among the highest, indicating a stronger deviation from what would be expected under the null hypothesis.

The variability in p-values and Z-scores across participants highlights differ-

Permutation Test Between Trials 1 and 12

Participant ID	Observed JSD	Permutation P-value	Z-score
001	0.8321	0.0228	2.2763
002	0.723	0.4831	0.7014
003	0.8305	0.0036	2.9147
004	0.7644	0.3933	0.8537
005	0.7864	0.1287	1.5191
006	0.8326	0.0226	2.2793
007	0.8016	0.3568	0.9215
008	0.8173	0.0641	1.8514
009	0.7972	0.6688	0.4277
010	0.8183	0.5271	0.6325
011	0.7709	0.4274	0.7936
012	0.8176	0.1939	1.2991
013	0.8186	0.2591	1.1285
014	0.7509	0.416	0.8134
015	0.8021	0.3161	1.0026
016	0.7057	0.5075	0.6627
017	0.8029	0.3199	0.9947
018	0.8298	0.0264	2.2195
019	0.8309	0.0269	2.2128
020	0.8272	0.03	2.1702
			Stouffer's P: 6.083e-10

Table 3.1: Permutation Test Between Trials 1 and 12.

ences in individual synergy uses across trials. The Stouffer's method, which aggregates p-values across all participants, yields a Stouffer's P-value of 6.083×10^{-10}, suggesting that, collectively, the results provide strong evidence of an effect when comparing Trials 1 and 12. This Stouffer's P indicates a consistent trend across participants, supporting the hypothesis that there is a difference in motor synergy behaviour use between the two trials across the cohort.

Results of the permutation test between Trials 13 and 36 in Table 3.2 show subtle variability across participants in the observed JSD values and corresponding p-values. Notably, Participant 009 has a lower permutation p-value (0.0112), suggesting a significant deviation when compared to other participants, with a corresponding Z-score of 2.5355. This stands out in contrast to the generally higher p-values observed among most other participants, indicating a lack of significant divergence between the trials for the majority of the cohort.

Stouffer's combined p-value of 0.1065 suggests that, when considering all participants collectively, there is no strong evidence of significant differences between Trials 12 and 36. This aligns with the generally non-significant p-values for individual participants, implying that the transition between these trials does not exhibit

Permutation Test Between Trials 12 and 36

Participant ID	Observed JSD	Permutation P-value	Z-score
001	0.8055	0.3724	0.8921
002	0.6119	0.8466	0.1935
003	0.4169	0.985	0.0188
004	0.4043	0.9645	0.0445
005	0.4633	0.9721	0.035
006	0.5681	0.7679	0.2951
007	0.4554	0.931	0.0866
008	0.4066	0.989	0.0138
009	0.8326	0.0112	2.5355
010	0.6235	0.8024	0.2502
011	0.3905	0.9668	0.0416
012	0.6548	0.7542	0.3131
013	0.7429	0.3663	0.9034
014	0.6653	0.6777	0.4157
015	0.7945	0.4672	0.727
016	0.3753	0.9754	0.0308
017	0.4596	0.8104	0.2399
018	0.3812	0.9785	0.027
019	0.4742	0.9174	0.1037
020	0.4269	0.9591	0.0512
			Stouffer's P: 1.065e-01

Table 3.2: Permutation Test Between Trials 13 and 36.

a pronounced overall effect and no significant change in motor synergy use.

Permutation Test Between Trials 36, Day 1 and Trial 3, Day 2

Participant ID	Observed JSD	Permutation P-value	Z-score
001	0.3498	0.9868	0.0165
002	0.401	0.9683	0.0397
003	0.342	0.9973	0.0034
004	0.07	0.964	0.0451
005	0.0715	0.9685	0.0395
006	0.062	0.9693	0.0385
007	0.2415	0.9982	0.0022
008	0.4767	0.9278	0.0906
009	0.2129	0.9994	0.0008
010	0.2048	0.9849	0.019
			Stouffer's P: 9.256e-01

Table 3.3: Permutation Test Between Trials 36, Day 1 and Trial 3, Day 2.

Permutation test results between Trials 36, Day 1, and Trial 3, Day 2 show little variation across participants, as evidenced by consistently high permutation p-values. The observed JSD values range from 0.070 to 0.4767, and the corresponding Z-scores are notably low, indicating a lack of strong deviation from the null hypothesis. The consistently high p-values (mostly above 0.96) suggest that there is no significant difference in the observed motor synergy distributions between these trials for the majority of participants. This consistency is underscored by the Stouf-

fer's p-value of 0.9256, which suggests that the differences across motor synergy distributions between the two trials are not statistically significant.

Overall, data indicate that the transition between Day 1 and 2 exhibits stability in participants' motor synergy use, as reflected in the low Z-scores and the high permutation p-values. This could imply that motor synergies following overnight sleep dependent consolidation and 'warm up' on the following day do not undergo significant alterations, suggesting a possible plateau for synergy development.

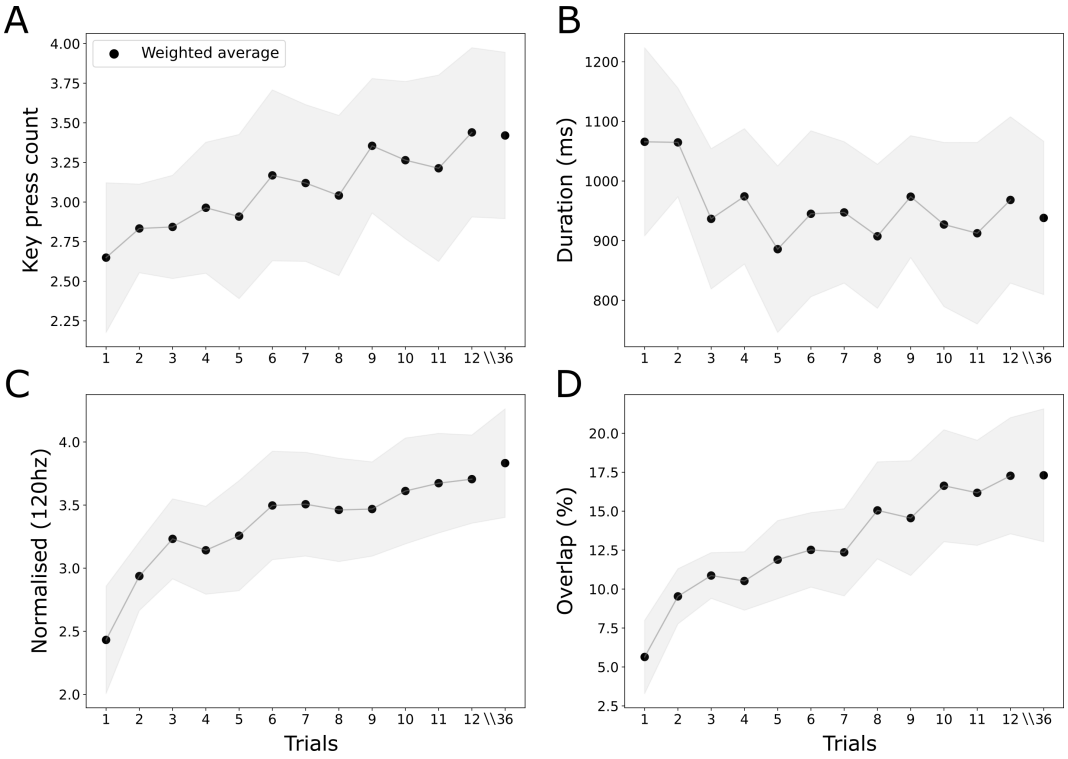

Figure 3.2: Kinematic Changes Over Early Learning. Weighted averages and confidence intervals are shown in each plot. Key press count (A) shows the average number of key presses made by participants in each synergy behaviour. Duration of synergy behaviour (B) shows the average duration of synergy behaviour in milliseconds. Normalised data (C) shows a standardised score of key presses made when key press counts are normalised to 120. This data shows how key press counts per synergy behaviour increase over trials regardless of duration. Overlap data (D) shows the extent to which, in percent of total frames of the synergy behaviour, digit movements "overlapped" with one another. Generally, overlap means that a proceeding key press movement (n+1) had a detectable onset prior to the offset of the preceding key press movement (n).

Kinematics data - including Key press count, Duration (in ms), Normalised Hz, and Overlap - over the course of early learning Trials 1 to 12, and up to Trial 36, are shown in Figure 3.2. A trend of improvement is observed across all metrics. Key

press count weighted average increases from 2.649 (Trial 1) to 3.440 (Trial 12), a gain of 0.791 ± 0.364 (SEM), indicating enhanced speed and increased consistency as denoted by the standard error of the mean. Duration, which is the weighted average of synergy durations used in a trial, shows a reduction from 1065.804 ms (Trial 1) to 968.286 ms (Trial 12), indicating a slight decrease in synergy lengths over early stage learning. The difference of -97.517 ± 107.387 suggests considerable variance in individual performance, as seen in the large SEM values. Keypress count normalised to 120 Hz also displays a positive trend, increasing from 2.432 to 3.705 between Trials 1 and 12, with a difference of 1.274 ± 0.281, reflecting that synergies contained more keypresses over early stage learning regardless of duration. This trend is further supported by the increase in overlap percentage, which rose from 5.642% in Trial 1 to 17.273% in Trial 12, with a difference of 11.631 ± 2.249. The increase in overlap suggests that participants became better at chunking individuated keypress movements, thereby reducing the temporal gaps between successive actions. Overall, the observed changes point to an improvement in both temporal and spatial integration of motor actions, consistent with the development of multi-effector synergy as participants progressed through early stage learning.

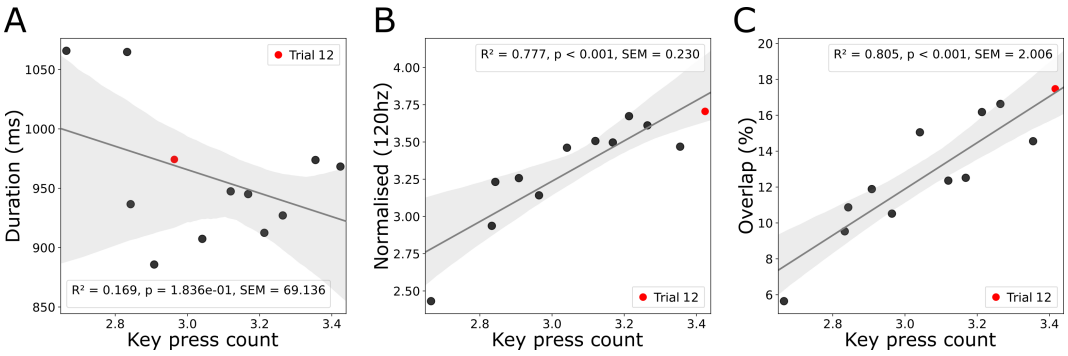

Figure 3.3: Each subplot displays a linear regression (gray line) with shaded 95% confidence intervals, illustrating the relationship between mean keypress count and three kinematic measures: (A) Duration of synergy behaviour (ms), (B) Normalised keypress count (per 120 frames), and (C) Overlap percentage. Each dot represents one trial in early learning (12 trials shown). Inset statistics include the coefficient of determination R^2 (explaining the proportion of variance accounted for by the linear model), corresponding p-value, and standard error of the mean (SEM).

Fig 3.3 shows three correlations with original Keypress count, including Duration in ms, Normalised key press count in Hz, and Overlap in %. The relationship

between key press count and Duration (in ms) revealed a weak negative correlation ($R^2 = 0.169$, $p = 0.1836$), suggesting that increased key presses are loosely associated with a decrease in synergy duration, albeit not significantly. This could imply that the length of motor synergies does not necessarily decrease as a function of increasing keypress count per synergy over early stage learning. The second and third plots examine the associations between Key press count and two additional metrics - normalised frequency (120 Hz) and overlap percentage - both of which demonstrate strong positive correlations ($R^2 = 0.777$ and $R^2 = 0.805$, respectively, and $p < 0.001$). These correlations indicate that as participants increase their number of key presses, they also show a greater number of accurate keypresses per 120 frames (as seen in normalised frequency) and greater overlap between digit movements, reflecting the emergence of motor synergies. This pattern shows how skill in motor sequences is characterised by increased speed, yes, but also enhanced consistency and overlap in the execution of multi-effector synergies. A summary of the study results is presented below in Table 3.4.

Comparison	**Measure**	**Result**
Trials 1–12 (Day 1)	Stouffer's P	6.083×10^{-10}
Trials 13–36 (Day 1)	Stouffer's P	0.1065
Trials 36 (Day 1) vs Trial 3 (Day 2)	Stouffer's P	0.9256
Keypress Count (Trial 1–12)	Mean Δ per synergy	$+0.791 \pm 0.364$
Synergy Duration (Trial 1–12)	Mean Δ in ms	-97.517 ± 107.387
Normalized Keypress (Trial 1–12)	Mean Δ per 120 Hz	$+1.274 \pm 0.281$
Overlap (Trial 1–12)	Mean Δ in %	$+11.631 \pm 2.249\%$

Table 3.4: Summary of key statistical results from Study 1, including permutation test results and kinematic changes. Stouffer's P values are derived from participant-level permutation tests comparing JSD values between trial distributions. Reported kinematic measures reflect the mean change (Δ) and standard error across early learning.

3.4 Discussion

Study 1 investigated the early stages of motor sequence learning by analysing the development of motor synergies as participants practised a five-element keypress

sequence over two training sessions. I characterised motor synergies at the trial level by identifying clusters in dimensionally reduced pose data using HDBSCAN, considering each cluster as a motor synergy. I examined and confirmed these synergies by calculating the underlying kinematics of each HDBSCAN-determined cluster. Motor synergy development over time was quantified by calculating JSD scores between synergy distributions across trials. My primary hypothesis was that significant motor synergy development would occur in the early trials, stabilising as participants reached a performance plateau. Results confirmed this hypothesis, as synergy metrics indicated rapid development within the first 12 trials, followed by minimal changes in later trials. Moreover, kinematic data suggested an improvement in multi-effector coordination, with increased keypress overlap and consistency across trials. Collectively, these findings show the timeline of motor synergy formation and its impact on motor sequence skill learning.

3.4.1 Timing of Motor Synergy Formation in Early Learning

The aim of Study 1 was to examine the timeline and process of motor synergy emergence during the early stages of learning a novel motor sequence skill. While previous research has explored motor synergies in various contexts, few studies have investigated their development at this granular timescale. By leveraging a time-resolved analysis of motor sequence learning, this study shows when and how motor synergies begin to form and stabilise.

Findings indicate that motor synergy formation begins at the start of learning, with significant adjustments observable across initial trials. The rapid emergence of synergies suggests that even in the earliest stages of learning, participants are capable of initiating coordinated movements that optimise task performance. These findings align with theories suggesting that the nervous system can quickly integrate sensory feedback and motor output to facilitate efficient task execution (d'Avella and Bizzi, 2005; d'Avella et al., 2006). By employing high-resolution kinematic data and using JSD to quantify changes in synergy distributions over time, I provide quantitative evidence of rapid and stable motor adjustments. This interpretation is further supported by the internal consistency of synergy labels over time. Because

clustering was applied in a reduced wavelet space, substantial drift in movement structure would have led to reassignment into a different cluster or designation as noise. Thus, the persistence of synergy labels across trials implies behavioural recurrence, not silent drift.

It is important, however, to contextualise these results within the specific task constraints of the study. The use of a fixed keypress sequence, which participants were instructed to perform as quickly and accurately as possible, inherently binds the complexity of motor learning involved, not only to the task but also to the requirements, i.e., speed. Such tasks, with their low skill ceiling, are well-suited for capturing rapid skill gains but may not fully reflect the processes involved in acquiring more complex motor skills, e.g., a tennis serve that involves the entirety of the body. The observed rapid synergy formation may be more reflective of the task's structured, repetitive nature rather than a universal characteristic of motor learning (Kitago and Krakauer, 2016).

Given this understanding, while the study suggests that synergies can develop quickly within motor sequences executed in a continuous, circular manner, it remains speculative to assume that similar patterns would be observed in more complex motor learning contexts, such as those involving more intricate requirements. As such, while these findings may inform strategies for optimising early-stage learning in highly structured tasks, their applicability to broader training or rehabilitation contexts should be approached critically.

In particular, the idea that targeted interventions could enhance synergy formation in early learning phases may hold for tasks with clearly defined motor patterns, but it remains uncertain whether such strategies would be effective in scenarios that have more variable motor patterns. Further research exploring a wider range of motor tasks, especially those with higher skill ceilings and more complex skill demands, is needed to determine the generalisability of these findings.

3.4.2 Notes on Participant Specific Analyses

Analysing motor synergy development at the participant level allowed for a detailed capture of each participant's unique learning trajectory. This individualised

approach highlighted the variability in synergy formation, revealing distinct kinematic patterns that might have been obscured by pooling data across participants. Indeed, the range of motor synergy use each participant exhibited in response to the task demands discovers a nuanced view of motor synergy development.

A limitation of this participant-specific approach was the challenge of embedding alignment across participants. Individual UMAP embeddings created for each participant could not be merged post hoc for direct group-level comparisons, limiting my ability to generalise findings across the cohort. This lack of a standardised embedding space meant that while I captured fine-grained details for each participant, any group-level similarities in synergy patterns remained less accessible.

Further, computational demands were high as each participant's data underwent independent processing. Although resource-intensive, this approach preserved the integrity of each participant's unique motor patterns, which was crucial for understanding variability in early motor learning. However, this individualised method would be impractical to scale in larger settings, such as hospitals, where streamlined, efficient processing across patients is essential. Implementing individualised analysis on a large scale would require significant computational resources and time, which may not be feasible in real-world clinical or remote care contexts.

In summary, individual participant analysis in Study 1 provided valuable insights into personalised trajectories of motor learning but also presented challenges in generalising these findings across the cohort. Addressing these limitations could refine future approaches to analysing motor synergy development and improve applicability in larger-scale settings.

3.4.3 Implications for Training and Rehabilitation

Findings from Study 1 highlight the value of focusing on early-stage learning of skill acquisition in healthy individuals. The rapid emergence of motor synergies observed in the data suggests that the initial phase of practice could be particularly conducive to fostering coordinated movement patterns (d'Avella et al., 2006). For athletes, musicians, and other skill-based performers, early interventions that emphasise multi-effector coordination could accelerate the integration of individual

actions into fluid, efficient sequences. Structuring practice drills to focus on specific key transitions may help reinforce the synergy structures that support seamless execution. Additionally, introducing task variations after foundational synergies are established could further enhance adaptability, enabling performers to adjust to new or complex motor challenges.

However, translating these findings to rehabilitation, especially for individuals with neuromotor impairments such as stroke or traumatic brain injury (TBI), necessitates a more cautious approach. Research indicates that such patients often experience disruptions to corticospinal pathways, leading to impaired motor control and the reliance on compensatory, rather than adaptive, movement patterns (Kitago and Krakauer, 2016). Unlike the rapid synergy formation seen in healthy participants, patients recovering from neurological damage may not follow the same trajectory due to structural and functional changes in the brain.

Instead of exclusively prioritising the rapid development of new motor synergies, rehabilitation protocols may benefit from a balanced approach that focuses on both optimising the functionality of existing, albeit compensatory, synergies while still attempting to promote the formation of new, more adaptive patterns. While some patients may show the potential for synergy reformation, others might reach a plateau, where further attempts to modify deeply ingrained compensatory strategies yield diminishing returns. For these individuals, the focus should shift towards refining existing motor patterns to improve motor efficiency, reduce biomechanical strain, and lower the energy costs of movement.

At the same time, insights gained from the observed timeline of synergy stabilisation in healthy individuals can still inform rehabilitation practices. Early, intensive practice sessions can be useful, but the emphasis should be on clinical monitoring to detect when a patient may have reached the limits of their capacity for synergy development. An adaptive strategy ensures that the focus remains on realistic and patient-centred rehabilitation goals, where training intensity and strategy can be adjusted based on observed progress.

Simultaneously encouraging the formation of new motor synergies and op-

timising existing, compensatory patterns may tailor rehabilitation programmes to meet the individual needs of patients. Continuous monitoring and reassessment would be essential to identify patient maximums for motor synergy modification. At that point, the rehabilitation strategy could shift towards optimising the efficiency and functionality of existing motor strategies to support sustained improvements in daily activities and quality of life.

Ultimately, a flexible and responsive approach - grounded in both the observed potential for early synergy formation in healthy individuals and the realities of neurological recovery - could lead to more effective, individualised interventions that maximise patient outcomes without setting unattainable expectations. Future research is needed to explore how this balanced strategy of promoting new synergies and optimising existing patterns can be systematically implemented in clinical settings to enhance motor recovery.

3.4.4 Limitations and Methodological Considerations

One important limitation of this study, and others like it (Bonstrup et al., 2019; Buch et al., 2021), was the brief training period, which consisted of only two practice sessions within a 24-hour span. While the study design intentionally targeted the early stages of motor learning, a longer training period may be necessary to observe the stabilisation and refinement of motor synergies beyond the initial rapid learning phase. The short duration likely emphasised the formation of synergies that support immediate performance gains rather than those that underpin long-term skill retention and adaptability.

Moreover, the task required participants to execute the sequence at maximum speed and accuracy from the outset, which may have influenced the nature of the synergies observed. In many domains, especially sports and high-level motor performance, the capacity to execute tasks rapidly and precisely is a hallmark of expertise. However, for participants acquiring a de novo skill, the speed requirement may have compelled participants to form motor synergies more rapidly and with less explicit consideration than they would under less demanding conditions.

In this context, the results emphasise how motor synergies might develop un-

der conditions that prioritise immediate performance over gradual refinement. Such conditions may encourage the formation of synergies that benefit high-speed execution but are not yet optimised for slight variations or the ability to adjust to new movement goals - characteristics that would emerge with extended practice under variable conditions. This limitation suggests that the synergies detected in this study could be specific to scenarios that emphasise rapid skill acquisition under pressure, which may differ from synergies formed through prolonged, gradual practice. Thus, while the study captures an intriguing aspect of early-stage motor learning, future research should explore how these initial synergies evolve over extended periods and whether they transition into more flexible patterns with continued practice.

These limitations point toward two promising directions for future research. First, longitudinal studies spanning weeks or months of practice would enable researchers to track whether early-emerging synergies become further refined, reorganised, or supplanted by more efficient patterns over time. Such work could clarify whether early-stage coordination structures form the basis for long-term motor habits or whether they reflect transient solutions tailored to initial task demands. Second, incorporating more complex or ecologically valid tasks - such as variable sequence structures, dual-task interference, or real-world motor challenges - could reveal how synergies adapt under pressure, generalise across contexts, or break down under cognitive load. Together, these extensions would help distinguish transient coordination from robust, transferable motor skills.

Chapter 4

MOTOR SYNERGY CHANGES LARGELY FOLLOWING PERIODS OF REST INTERLEAVED WITH PRACTICE

"Sweet is the rest after labour."

— Ovid

4.1 Introduction

As shown in Study 1, early stages of motor sequence learning are characterised by rapid performance improvements. During this phase, learners transition from executing individuated movements to performing multi-effector motor sequence synergies with proficiency (Karni et al., 1995). Understanding the fine-grained time course and mechanisms underlying these rapid changes in motor behaviour is crucial for developing effective training protocols and rehabilitation strategies.

Recent research has highlighted that a significant portion of early learning gains occurs not only during active practice (online), but also during short rest periods interleaved with practice (offline). These micro-offline gains - defined as

short-term performance improvements observed between successive practice blocks - suggest that consolidation processes may operate within a single training session, not solely between sessions (Bonstrup et al., 2019). This challenges the traditional view that continuous practice is essential for skill acquisition, and instead underscores the functional role of rest in the learning process.

This phenomenon raises important questions about the kinematic mechanisms that facilitate rapid skill consolidation during short rest periods. While it is established that rest can enhance performance, it remains unclear whether specific motor patterns or synergies develop during these brief intervals (Fricke et al., 2020). If rest contributes to the emergence of motor synergies, incorporating strategic breaks into training could optimise learning outcomes. Conversely, if synergies form mainly during active practice, rest may serve to stabilise rather than enhance these patterns.

Study 1 showed that motor sequence skill improvements in early learning reflect both increased movement speed and consistent overlap in digit movements. However, the precise timeline of these changes - whether they emerge during rest or practice - remains unknown. Some evidence suggests multi-digit synergies may appear rapidly, within seconds, while other findings point to a slower progression over minutes or hours (Gentner et al., 2010). Clarifying when and how these synergies develop will inform strategies for optimising motor learning and rehabilitation protocols (Dayan and Cohen, 2011).

Current Study

The present study aims to investigate the influence of rest periods on the development of motor synergies during the early stage of learning a novel motor sequence skill. Specifically, I examine whether significant synergy formation occurs during rest periods (offline) or active practice periods (online) within training sessions. Twenty healthy participants practised a five-element keypress sequence over 36 trials, with practice sessions interleaved with rest periods. The methods for data collection were identical to those used in Study 1 on synergy development across early and late learning stages.

I hypothesized that motor synergies would develop primarily during rest pe-

riods within early learning trials, evidenced by significant measures of Jensen-Shannon divergence in motor synergy distributions during these intervals. This would suggest that rest facilitates rapid consolidation and refinement of motor synergy executions, contributing to performance improvements upon resumption of practice. An alternative hypothesis considered that motor synergies may primarily form during active practice, with rest periods contributing to stabilization rather than further development.

To quantify synergy development, I analysed changes in motor behaviour, i.e. distributions of synergy use over trials, using JSD. Skill gains were assessed by measuring the number of correct sequences performed per second. By comparing synergy development and performance improvements during rest and practice periods, I aimed to show the temporal dynamics of motor learning and the specific role of rest in facilitating skill acquisition.

Significance and Implications

Understanding how periods of rest influence the development of motor synergies during early learning has implications for neuroscience, psychology, physical therapy, and skill training. If rest periods substantially contribute to synergy formation, incorporating strategically timed breaks into training programs could enhance learning. This approach may be especially beneficial in rehabilitation settings, where patients need to relearn motor skills and integrate them into functional activities.

By identifying the conditions that facilitate rapid skill consolidation, practitioners can develop strategies to optimize motor learning outcomes. Recognizing the critical role of rest periods allows for tailored interventions that promote the development of efficient movement patterns.

This research advances our understanding of early motor skill acquisition mechanisms. By elucidating the impact of rest periods on motor synergy development, this work provides valuable insights that can optimize training protocols across various applications, ultimately improving outcomes for individuals seeking to enhance or restore motor function.

4.2 Methods

Methods for Study 2 are exhaustively detailed in Chapter 2, but I will provide a brief overview here of those elements most relevant to the proceedings.

4.2.1 Overview

In Study 2, I assessed motor sequence skill changes over specific periods of practice and rest during early-stage learning. Twenty participants learned the same motor sequence task as in Study 1 - typing the sequence 4-1-3-2-4 with each digit specifically assigned to one number: the index digit for number 4, the middle digit for number 3, the ring digit for number 2, and the little digit for number 1. Performance was measured as correct sequences per second (cs/s) across all trials. For each participant, the trial at which 95% of total skill gains were achieved was identified and denoted as T. This marked the end of early-stage learning and defined the window over which subsequent behavioural and kinematic changes were analysed. This threshold was chosen to isolate the rapid initial phase of learning, following prior work showing that early trials are most sensitive to micro-timescale learning dynamics (Bonstrup et al., 2019). Micro-online learning, occurring during practice periods, was defined as the change in cs/s from the beginning to the end of a trial - specifically, the difference between the first and final full seconds. Micro-offline learning, corresponding to rest periods, was defined as the change in cs/s between the last second of one trial and the first second of the next. These one-second windows were selected because the cs/s metric is computed at a one-second resolution, making it the most appropriate timescale for capturing fine-grained fluctuations in performance. To assess changes in motor synergy use during these same intervals, I computed the Jensen–Shannon Divergence (JSD) between synergy label distributions at the beginning and end of each micro-online and micro-offline period. Cumulative JSD values over the first T trials were used to quantify total synergy change during practice and rest. Digit movements were extracted using DeepLabCut (DLC), transformed via Morlet wavelets, and embedded into a two-dimensional UMAP space. HDBSCAN clustering was then used to segment behavioural synergies from the resulting time series. Unlike in Study 1, where permutation tests were

used to compare specific trial pairs, the goal here was to assess whether the magnitude of micro-timescale synergy changes predicted overall learning. Accordingly, I fit two linear regressions: one testing whether mean micro-online JSD predicted total cs/s gain over early learning, and another for mean micro-offline JSD. This approach allowed us to directly examine whether within-trial or between-trial synergy reorganisation better accounted for individual differences in learning trajectory.

4.2.2 Inclusion criteria

Participants were recruited if they were between 18 and 90 years old, spoke English, and were right-handed. Additional inclusion criteria required participants to have an internet-connected home computer with a USB port or USB port adapter, be able to participate in a videoconference with a member of the research team, and be capable of typing with their left hand without pain or discomfort. Participants had to be located within the United States to receive study materials. Exclusion criteria included being a staff member of Human Cortical Physiology Section (HCPS; my section with the NIH). If participants were unable to receive materials within the United States, they were required to withdraw from the study or defer participation until they could receive materials in the country.

4.2.3 Participants

Twenty naive right-handed healthy participants (16 women; mean± SD age 22.4 ± 2.63) gave their written informed consent to participate in the project, which was approved by the Combined Neuroscience Institutional Review Board of the National Institutes of Health (NIH). Active musicians were excluded from the study(Abraham and Drory, 2014). The sample size was determined via a power analysis of skill learning data collected using the same task (Censor et al., 2014).

4.2.4 Study design

Study 2 used a within-subject, repeated-measures experimental design to investigate changes in motor synergy during motor sequence learning over short-term practice sessions, focusing specifically on micro-online and micro-offline assessments within early-stage learning. Participants learned the same motor sequence task -

typing the sequence 4-1-3-2-4 as quickly and accurately as possible, with each digit specifically assigned to one number. Over a single day, each participant completed 36 trials, with each trial consisting of 10 seconds of practice followed by 10 seconds of rest, ensuring consistent practice exposure across individuals. Outcome measures included cumulative cs/s and JSD values during micro-online (practice) and micro-offline (rest) periods within early stage of learning.

A within-subject design with repeated measurements allowed for direct comparison of each participant's performance and motor synergy use during the micro-on and -offline periods. I assessed whether performance gains and changes in motor synergy occurred mainly during active practice or rest intervals within early-stage learning. The design minimized inter-subject variability and enhanced statistical power, making it well-suited for detecting significant changes in motor synergy use and its relationship to micro-level learning processes during skill acquisition.

4.3 Results

Figure 4.1 illustrates skill learning from Study 2, depicting the progression in sequence speed (cs/s) across trials. The black line shows the trial means, while the red line represents the fitted logarithmic model. The shaded gray area around the trial means indicates the standard error of the mean (SEM), capturing variability in performance across participants.

At Trial 1, mean cs/s is 0.444 (SEM ±0.034), establishing the baseline skill level. By Trial 11, where participants reach 95% of their skill gain, mean cs/s has increased to 0.853 (SEM ±0.043). The dashed line marks the point of 95% saturation in learning. The total gain from Trial 1 to Trial 11 is 0.409 cs/s (SEM ±0.055), highlighting the rapid improvement in early trials. The close alignment of the trial means with the logarithmic model shows the model's accuracy in capturing the early learning trajectory.

Figure 4.2 presents cumulative cs/s gains over the first 11 trials in Study 2, segmented into micro-online and micro-offline periods. Each data point reflects the cumulative change in correct sequence speed, with red representing online gains

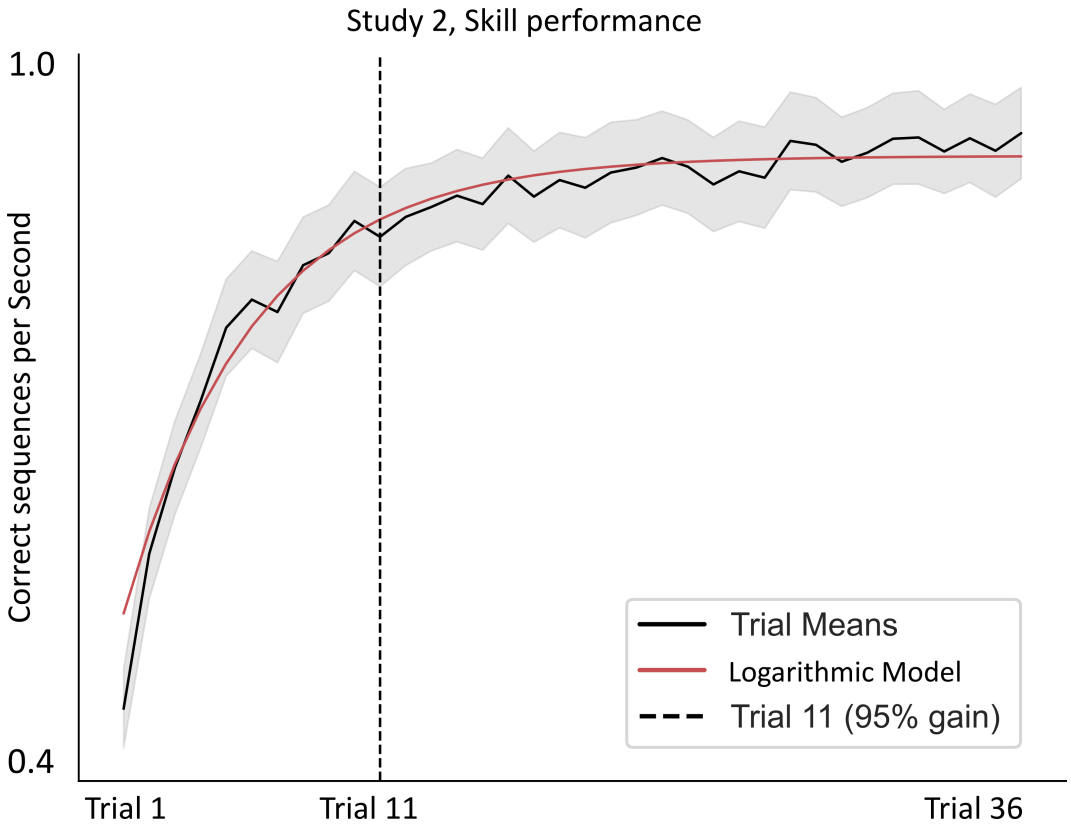

Figure 4.1: Study 2, Skill Performance (N = 20). Skill is measured as Correct Sequences per Second (cs/s). The black line shows the mean performance across participants, with the shaded region representing ±1 SEM. The red curve is a fitted logarithmic model capturing early learning. The vertical dashed line marks Trial 11, which corresponds to the point of 95% total skill gain (T), defining the end of early-stage learning.

over practice and blue representing offline gains following rest. The y-axis, labelled as cumulative cs/s, captures the progressive skill learning across trials.

Data highlights that offline gains, or the improvements occurring after rest, consistently surpass online gains, underscoring that the majority of early learning can be attributed to offline changes. For instance, at Trial 1, the cumulative offline change reaches 0.187 cs/s (SEM ±0.028), while the online gain is 0.063 cs/s (SEM ±0.034). This pattern persists through Trial 11, where cumulative offline gains continue to drive the overall learning trajectory, suggesting that the offline periods play a substantial role in skill acquisition during the early learning phase. Notably, between the last offline gain at Trial 10 (0.529 cs/s, SEM ±0.059) and the final online gain at Trial 11 (0.409 cs/s, SEM ±0.055), there is a slight reduction, indicating a stabilization as participants approach the 95% gain mark. The SEM bars indi-

cate variability across participants, but the trend consistently favours offline gains, demonstrating the significance of rest in skill consolidation.

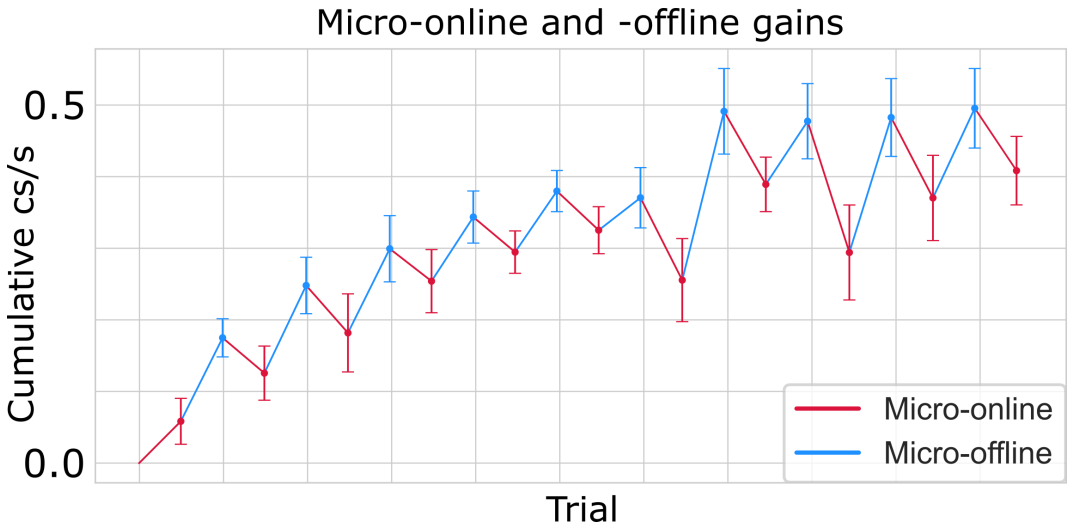

Figure 4.2: Cumulative correct sequence speed (cs/s) gains are plotted across the first 11 trials of Study 2 (N = 20). Blue represents micro-offline gains (rest periods), and red represents micro-online gains (practice periods), with each point showing group means ±1 SEM. Offline gains consistently exceed online gains across trials, suggesting that brief rest periods contribute substantially to early learning.

Figure 4.3 illustrates cumulative sequence speed (cs/s) gains and motor synergy use over early learning in Study 2, with each dot representing an individual participant's mean. The left panel shows cumulative skill change segmented into total, micro-online, and micro-offline gains. Total early learning is 0.409 cs/s (SEM ±0.055), with offline gains of 1.406 cs/s (SEM ±0.317) markedly outweighing the online losses of -0.997 cs/s (SEM ±0.325). This pattern shows the dominant role of offline, rest-driven gains in early learning, as offline periods account for the majority of cumulative skill acquisition.

The right panel presents JSD values for the change in motor synergy use through early learning. The offline JSD mean is 0.596 (SEM ±0.026), substantially higher than the online mean of 0.316 (SEM ±0.028). These values suggest that motor synergy use largely changes following offline intervals, paralleling the cumulative skill gains. A paired t-test confirms a highly significant difference between online and offline JSD ($p < 2.086e\text{-}15$, $t_{19} = -23.22$), reinforcing the conclusion that offline periods are critical for consolidating motor synergy adjustments

essential to skill learning.

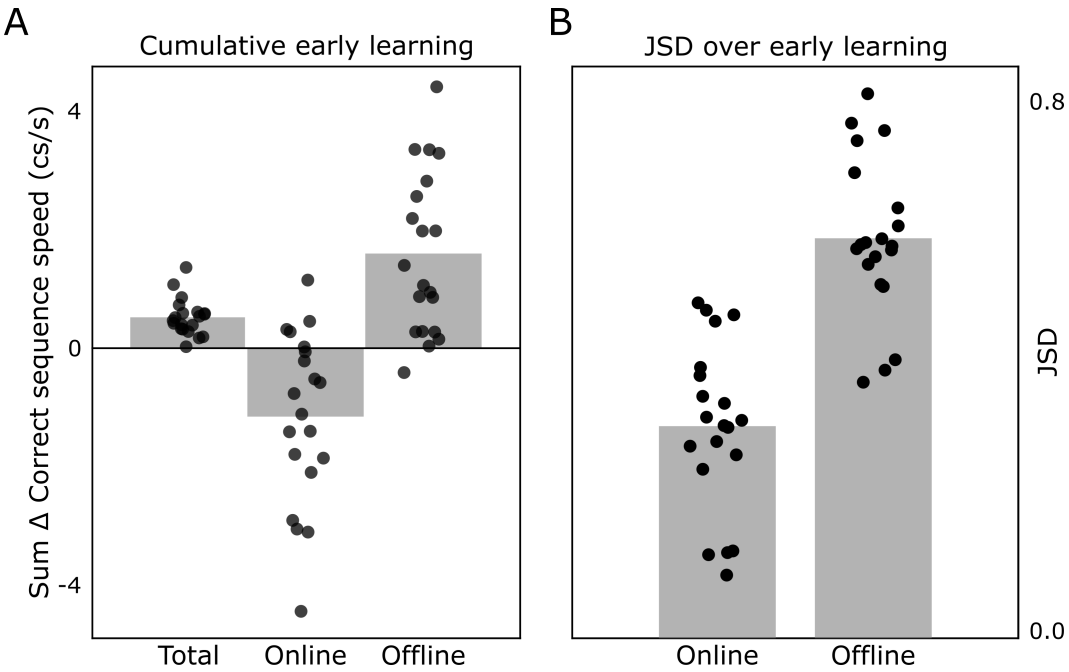

Figure 4.3: Skill and Synergy Summary Over Early Learning.

Figure 4.4 shows the relationship between motor synergy use – over offline and online periods – in JSD, and total early learning gains (cs/s) in Study 2. The left panel shows a significant positive correlation between offline JSD and total early cs/s gain ($R^2 = 0.7112$, $p < 0.001$, SEM = 0.2924). This strong relationship suggests that changes in motor synergy use occurs following offline periods of rest and are closely associated with improvements in total skill learning. In contrast, the right panel illustrates the relationship between online JSD and total early learning gains, where no significant correlation is observed ($R^2 = 0.1037$, $p = 0.1661$, SEM = 0.4850). The lack of association between online motor synergy use and early learning gains demonstrates the importance of offline processes in facilitating skill learning, indicating that offline synergy use change is likely integral to early stage skill learning and consolidation.

4.4 Discussion

Study 2 examined the temporal dynamics of motor synergy development during early-stage motor sequence learning, focusing on the role of alternating practice

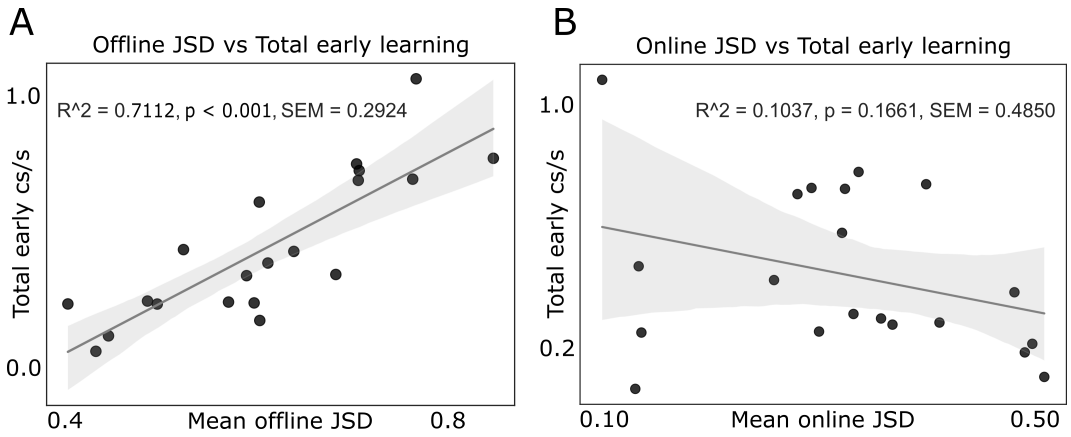

Figure 4.4: Offline and Online JSD Correlations with Total Early Learning.

and rest periods. Training to learn a new skill commonly involves practice of alternating periods of practice and rest (Christiansen et al., 2018; Song and Cohen, 2014; Maier et al., 2019). Recent work showed that substantial performance gains during early learning develop during rest periods interspersed with practice, a form of rapid consolidation of skill (Bonstrup et al., 2019). In that vein, little is known about the morphology of synergies over those periods of skill acquisition. To address those questions in this study, participants practised a five-element keypress sequence over 36 trials, with skill gains assessed by correct sequences per second (cs/s) and synergy development measured using Jensen-Shannon divergence (JSD) in motor behaviour distributions. My primary hypothesis was that offline rest periods would facilitate rapid motor synergy formation, contributing significantly to early learning gains. Results confirmed this hypothesis, with offline gains surpassing those during practice, as indicated by higher cumulative JSD values and significant correlations with total skill improvements. Offline periods not only enhanced performance but also promoted the transformation of motor synergies, characterised by increasingly efficient and overlapping digit movements. In the context of learned skills, multi-effector synergies may represent the manifestation of an optimal or most efficient modular control strategy used by the central nervous system when delicate manual coordination is required (Poggio and Bizzi, 2004; Gentner et al., 2010; Tresch and Jarc, 2009; Santello et al., 2016; Fricke et al., 2020). These results underscore the critical role of rest in skill consolidation, revealing a mechanism through which the

central nervous system optimises motor sequence execution during early learning.

4.4.1 Temporal Dynamics of Micro-Offline Learning

Study 2 replicates previous findings showing that the most substantial performance improvements in early motor sequence learning occur not during active execution but during rest periods interleaved with practice (Bonstrup et al., 2019). These micro-offline gains refer to short-term performance improvements observed between practice trials, and are quantified as increases in correct sequences per second (cs/s) from the end of one trial to the start of the next.

This phenomenon challenges the traditional paradigm that views rest as merely a break to prevent fatigue. Instead, these findings indicate that rest periods act as active processing windows, during which the nervous system consolidates recent motor experiences into more stable coordination patterns. One might think of rest periods as akin to a "compression algorithm" applied to recent motor sequences: the brain summarises, streamlines, and stabilises relevant patterns, reducing noise and improving efficiency on the next attempt.

The rapid onset of consolidation during these brief rest intervals suggests that motor memory is not only formed but also meaningfully reorganised in seconds - not hours - after practice. This aligns with theories of fast synaptic plasticity, where memory traces are strengthened immediately post-experience through reactivation and integration into ongoing motor representations (Albouy et al., 2013).

Recognising rest as an active contributor to skill acquisition has clear practical implications. Training protocols could be optimised by deliberately incorporating short, strategically timed breaks to capitalise on this consolidation effect. This is especially relevant in time-constrained settings such as rehabilitation or elite training environments, where efficient learning is essential. Future studies should further investigate how the timing and duration of rest intervals modulate consolidation to enhance motor performance.

4.4.2 Neural Mechanisms of Rapid Consolidation

The enhancement of motor skills during rest intervals observed in Study 2 highlights the critical role of neural mechanisms underlying rapid consolidation. Rest intervals of approximately 10 seconds, as used in this and previous studies (Bonstrup et al., 2019; Buch et al., 2021), have been shown to be sufficient for triggering measurable consolidation effects, suggesting that neural replay and plasticity mechanisms can initiate within seconds of rest onset. This reactivation is thought to strengthen synaptic connections, thereby stabilising newly acquired motor memories. Functional neuroimaging studies support this mechanism, showing increased activity in motor regions during rest following motor learning tasks (Albouy et al., 2013). Such activity suggests that the brain continues to process and refine motor information even in the absence of overt movement. Additionally, electrophysiological studies have demonstrated that synaptic plasticity markers, such as long-term potentiation, are enhanced during rest periods (Dong et al., 2021). These findings indicate a continuum of overlapping processes between practice and rest, rather than distinct phases of memory encoding and consolidation. The neural substrates engaged during rest may involve not only the motor cortex but also the hippocampus and other areas associated with memory processing (Tamaki et al., 2013). This interconnected activity underscores the complexity of motor learning and the importance of considering both practice and rest in training paradigms. Understanding the neural mechanisms of rapid consolidation has implications for optimising learning strategies. By tailoring practice schedules to align with periods of heightened neural plasticity during rest, it may be possible to enhance skill acquisition.

4.4.3 Motor Synergy Development During Rest Periods

Study 2 demonstrates that motor synergies evolve not only during active practice but also significantly during rest periods in early motor sequence learning. These offline intervals were associated with greater cumulative changes in motor behaviour distributions, as quantified by JSD, indicating substantial reorganisation of synergy structure compared to practice periods.

During rest, participants exhibited increased overlap and coordination among

digit movements, suggesting that the central nervous system uses these intervals to refine the temporal and spatial coupling of effectors. This likely reflects the consolidation of more efficient multi-effector synergies that support smoother, more integrated execution of the keypress sequence. These findings support the view that motor synergies are adaptable control modules that are reshaped through both experience and consolidation (Tresch and Jarc, 2009; Santello et al., 2016).

I propose that during short rest periods, the CNS may fine-tune inter-digit timing relationships to reduce temporal gaps and redundant motor output, optimising the modular organisation of motor commands. This refinement process may improve the efficiency of motor control, supporting more fluid and economical sequence production (Gentner et al., 2010; Poggio and Bizzi, 2004).

These results extend prior work by showing that not only does behavioural performance improve during rest, but the underlying structure of motor coordination also undergoes meaningful transformation - underscoring the functional importance of rest in shaping motor skill development.

4.4.4 Distinctions Between Online and Offline Learning

My findings reinforce findings that a clear distinction exists between online gains occurring during active practice and offline gains emerging during rest periods. While practice is essential for encoding new movement patterns, the disproportionate improvements observed during rest suggest that consolidation processes play a more substantial role in early skill development than previously recognised. This challenges the conventional emphasis on continuous practice for motor learning.

The differentiation between online and offline learning reflects distinct but complementary mechanisms of memory stabilisation and refinement. Online learning engages neural circuits associated with motor execution and immediate feedback processing (Classen et al., 1998). In contrast, offline learning during rest may involve synaptic reorganisation and the strengthening of neural pathways without the interference of new motor inputs (Shadmehr and Mussa-Ivaldi, 2012).

These distinctions have practical implications for task design and training efficiency. The assumption that skill improvements occur uniformly across practice is

not supported by the data. Instead, incorporating rest intervals strategically within practice sessions may maximise learning by allowing offline consolidation to occur. This approach could lead to more efficient use of training time and potentially reduce the total amount of practice required to achieve proficiency.

Further research is needed to explore the optimal balance between practice and rest. Variables such as the timing, duration, and frequency of rest periods may differentially affect learning outcomes. Additionally, individual differences in response to online and offline learning phases should be considered, as they may influence the effectiveness of training interventions.

Chapter 5

MOTOR SYNERGY DOES NOT TRANSFER BETWEEN SUBSEQUENTLY LEARNED SKILLS

"Oh, East is East, and West is West, and never the twain shall meet."

— Rudyard Kipling

5.1 Introduction

Motor sequence learning and the ability to transfer or generalise learned skills to new contexts are fundamental aspects of human behaviour. The nervous system's capacity to generalise learned motor patterns facilitates adaptation to novel environments and tasks, which is crucial for daily functioning (Adams, 1987; Schmidt, 1975). Understanding the mechanisms underlying motor skill transfer has significant implications for neuroscience, psychology, rehabilitation, and skill training (Gentile, 2000; Krakauer and Mazzoni, 2019).

Motor skill transfer is represented as the gain (or loss) in proficiency in one motor task as a result of practice on another task. Traditional research in motor

adaptation has extensively explored this phenomenon using controlled laboratory tasks, such as targeted reaching movements, where altered sensory feedback or dynamic conditions led to measurable performance errors that diminish with practice (Shadmehr and Wise, 2005; Krakauer and Shadmehr, 2006).

Evidence from adaptation studies suggests that transfer is frequently partial and highly dependent on the degree of similarity between the trained and transfer tasks (Krakauer, 2009; Taylor and Ivry, 2013; Taylor et al., 2014). For instance, adaptation to a visuomotor rotation during reaching movements can facilitate subsequent adaptation to different degrees of rotation, but the transfer is often incomplete and context-specific. Transfer tends to occur when tasks share common spatial, temporal, and biomechanical characteristics, involving similar workspaces, muscle groups, or movement patterns (Seidler et al., 2010; Wang and Sainburg, 2007).

Contrastingly, clinical observations provide compelling evidence of significant transfer between distinct, naturalistic tasks. Patients undergoing rehabilitation for motor impairments often show improvements in untrained tasks following targeted therapy (Lang et al., 2013; Wolf et al., 2006). For example, practising activities like writing or using utensils can lead to enhanced performance on standardised functional assessments, despite the tests not being directly practised during therapy. This suggests that motor skill transfer in real-world contexts may be more extensive than observed in controlled experiments. The discrepancy between laboratory findings and clinical observations raises important questions about the factors that facilitate or constrain motor skill transfer.

Specific transfer occurs when there is a close alignment between pre-existing coordination patterns and the demands of the new task, allowing for the direct application of learned skills. General transfer involves applying general capacities, such as perceptual skills or postural control, to new tasks where specific movement patterns may differ (Schmidt, 1975; Adams, 1987). Exploratory behaviours and movement variability play crucial roles in facilitating both specific and general transfer by allowing performers to adapt and discover functional movement solutions in novel environments (Latash et al., 2007). While brief, this distinction

underlies my approach to interpreting transfer and is summarised in Table 5.1.

Transfer Type	Definition	Example / Implication
Specific Transfer	Occurs when previously learned coordination patterns apply directly to a new task due to structural similarity.	Using a trained 3-element sequence (e.g., 4-1-3) within a 5-element sequence (e.g., 2-4-1-3-4). The 3-element synergy maps cleanly onto a portion of the new task.
General Transfer	Involves the application of broader perceptual-motor strategies or skills to novel tasks with different structures.	Improved postural control, attention, or timing flexibility from training one sequence aids a different, structurally unrelated task.

Table 5.1: Summary of specific vs. general transfer of motor sequences.

Motor skill transfer is not only influenced by the similarity between tasks but also by the amount and structure of practice (Wulf, 2008; Wulf and Lewthwaite, 2010). Early in learning, motor synergies are more flexible and can be adapted to new tasks. With increased practice, synergies may become more specialised and less transferable due to the consolidation of task-specific coordination patterns (Santello et al., 2016; Latash, 2021). This suggests an optimal window during which practice facilitates transfer before motor synergies become too rigid.

Furthermore, interference with motor memory consolidation during the immediate post-acquisition period can degrade not only the learning but also the transfer of motor skills to new tasks, underscoring the critical role of consolidation in facilitating skill generalisation (Robertson, 2019; Censor et al., 2014). The neural substrates for motor memory consolidation, which are essential for both retention and successful transfer, depend on the structure of practice. Interference to the primary motor cortex affects retention and transfer of the practised skill following constant practice but not variable practice, whereas interference to the dorsolateral prefrontal cortex affects retention and transfer following variable practice but not constant practice (Kantak et al., 2010; Krakauer and Mazzoni, 2019). When practice is organised using a variable structure, it enhances transfer performance on novel versions of the task more effectively than constant practice. Variable practice allows learners to compare and contrast different task versions, leading to deeper processing and the learning of abstract rules that facilitate both retention and trans-

fer to new skills (Wulf, 2008; Schmidt, 1975).

Statistical models of motor learning further suggest that the history of prior motor behaviour and context significantly influence transfer and interference between skills (Donchin et al., 2003; Wang and Sainburg, 2005). Generalisation depends not only on the similarity of contextual cues but also on the uncertainty associated with parameter estimates derived from previous experiences. The extent of transfer is thus modulated by the learner's history of interactions with different tasks and contexts, shaping their expectations and error generalisation processes (Kording and Wolpert, 2004; Gonzalez Castro et al., 2014).

Recent studies have investigated the role of context in motor learning, skill transfer, and the ability to learn and recall overlapping motor mappings (Wang and Sainburg, 2005; Taylor et al., 2014). For example, switching the effector used to perform a task can serve as a contextual cue that reduces interference and allows independent learning of opposite visuomotor rotations (Krakauer and Shadmehr, 2006). Contextual differentiation can facilitate the formation of separate motor memories, enhancing the ability to transfer skills without detrimental interference.

Despite extensive research on motor skill transfer, there remains a gap in understanding how motor synergies developed during the learning of simple tasks transfer to more complex tasks, particularly within the context of motor sequences. Previous studies have largely focused on adaptation in single, isolated movements or on tasks involving continuous dynamics, often neglecting the discrete and hierarchical nature of many real-world motor skills (Schmidt, 1975; Santello, 2002).

Current Study

The present study aims to address how motor synergies developed during the learning of simple tasks transfer to complex tasks, specifically within the context of motor sequences. While previous research has primarily focused on isolated movements or continuous tasks, this study investigates the discrete and hierarchical nature of motor sequences, which more closely resemble real-world motor skills. The study examines whether synergies formed during the practice of a three-element motor sequence can facilitate performance in a paired five-element sequence that

embeds the initial sequence. By manipulating the amount of practice on the initial sequence (3, 7, or 11 trials), the aim is to determine how practice duration impacts the flexibility and effectiveness of motor synergy transfer.

Study 3 hypothesises that synergies are more flexible during early stages of learning and that brief, varied practice may enhance transfer by preserving this flexibility. The study's design allows us to explore whether longer practice durations lead to more rigid, task-specific synergies that may limit transfer. The choice of three- and five-element sequences reflects the hierarchical nature of motor learning, where simpler learned components often underpin more complex behaviours.

The novelty of study 3 is in its focus on continuous motor sequences and the temporal aspects of synergy formation and transfer, which remain under-explored. By investigating how practice structure and consolidation processes influence the development and transfer of motor synergies, this study seeks to fill critical gaps in our understanding of motor skill learning and generalisation.

Significance and Implications

Despite extensive research on motor skill acquisition and transfer, significant gaps remain in understanding how motor synergies developed during the learning of simple tasks transfer to more complex sequences. Previous studies have primarily focused on adaptation in single, isolated movements or continuous tasks, often neglecting the discrete, hierarchical nature of many real-world motor skills. Furthermore, whilst the role of practice in enhancing motor skill retention is well-documented, less is known about how the structure and duration of practice influence the transfer of motor synergies. This gap is particularly pronounced when considering the temporal dynamics of motor synergy development and their applicability across tasks of differing complexity.

By addressing these gaps, this study has implications for both theoretical and applied domains. If brief early practice enhances transfer by preserving the flexibility of synergy structures, training programmes could emphasise varied and shorter practice sessions to optimise learning outcomes. This could be advantageous in rehabilitation settings, where patients recovering from motor impairments must re-

learn foundational motor skills and apply them to diverse, daily activities.

Identifying the specific conditions that facilitate synergy transfer can also guide strategies to minimise interference between tasks and enhance the consolidation of beneficial movement patterns. By considering factors such as practice duration, task complexity, and individual variability in learning dynamics, practitioners can develop tailored interventions that support both the acquisition and transfer of motor skills. Ultimately, this research contributes to a more nuanced understanding of motor learning and transfer, offering insights that could improve the design of training and rehabilitation protocols and promote more effective skill generalisation in both clinical and everyday contexts.

5.2 Methods

Methods for Study 3 differ from the exhaustive description set forth in Chapter 2. This section will begin with an overview, followed by one section for each method, technology, or analysis new to Study 3.

5.2.1 Overview

Study 3 examined motor sequence learning transfer by analysing correct keypresses per second (ckp/s), keypress transition speeds, and motor synergy at the trial level. Sixty healthy participants were randomly assigned to one of six orders of motor sequences, referred to as rotations, designed to reduce ordering effects and diversify sequence exposure. In each rotation, participants trained on a unique 3-element sequence (e.g., 4-1-3) for either 3, 7, or 11 trials, followed by a 5-element paired sequence (e.g., 2-4-1-3-4) in which the initial 3-element sequence was embedded, for the same number of trials.

The equal number of trials for each sequence pair allowed me to examine how synergies formed during the 3-element sequence transferred to the more complex 5-element sequence, under varying durations of practice. To preserve the ability to detect transfer specifically (rather than co-learning), I avoided including a warm-up period at the onset of the 5-element sequence. Unlike Study 1, where synergies were compared across trials of the same sequence, Study 3 required the immediate

comparison of a learned and a novel sequence to isolate transfer.

I used ckp/s to quantify skill learning (i.e., participants' speed and accuracy) through each trial. Keypress transition speeds, measured in milliseconds, served as an initial metric for evaluating skill transfer, specifically to assess whether the skill learned from the 3-element sequence transferred to the same embedded subsequence within the 5-element task.

To investigate group-level patterns, UMAP with Hamming distance was applied to all participant data within each rotation, providing a loose grouping of pose data related to specific 3-element sequences and their paired 5-element sequences. Motor synergies were identified at the trial level using recursive HDBSCAN, which segmented clusters iteratively to reveal synergy patterns relevant to transfer.

To assess whether motor synergies developed in the 3-element sequence transferred to the paired 5-element sequence, I performed JSD analyses between the final trial of the 3-element sequence and the first trial of its paired 5-element sequence (i.e., whether synergies formed during training were immediately deployed in the more complex context before any additional learning could take place). Permutation tests on JSD values determined empirical p-values, allowing us to statistically quantify the transfer of motor synergies. This approach provided insight into how motor synergies develop and transfer as participants progressed from simple to complex motor sequences under different training durations.

5.2.2 Apparatus

The data collection apparatus for Study 3 comprised several key components, each selected for their specific capabilities in capturing and recording participants' hand movements during motor sequences.

The frame of the data collection apparatus was constructed using 14 pieces of 45 mm x 45 mm T-slot aluminium, configured into a structure measuring 75 cm wide, 60 cm tall, and 60 cm deep. T-slot aluminium is a robust, modular building material widely used in research and industrial applications due to its adaptability and ease of assembly. This design provided a customizable framework, allowing precise placement of key components. The use of T-slot aluminium was advanta-

geous for securely affixing all components at optimal positions for data collection, while also allowing adjustments as needed for different experimental setups.

Figure 5.1: Study 3 apparatus and components. (Left) shows the full data collection frame; (Top-middle) Storm Interface 4-Key Polymer Keypad, (Top-right) Ultimarc U-HID Control Interface, (Lower-middle) SMALLRIG 5 in Magic Arm with Clamp, and (Bottom-right) GoPro HERO10 Black Camera used to record hand motion.

Figure 5.1 shows the completed apparatus for data collection, which includes all mounted components, each affixed to the frame in configurations facilitating data collection and accessibility. The modular T-slot design allows for the integration of future components if needed. I'll now offer a brief introduction to each of the components affixed to the apparatus.

The **Storm Interface 4-Key Polymer Keypad** input device featured four coloured buttons, each with 1.4 mm key travel and an actuation force of 150 g. The button box is ideal for consistent tactile feedback in experimental settings, and the high responsiveness is suited to capturing data for motor sequence keypress tasks with millisecond resolution.

The **Ultimarc U-HID Control Interface** features 50 connections that can be configured for various inputs, including pushbuttons, joysticks, and rotary encoders. Using the U-Config software and the U-HID SDK package, I assigned the numbers 1, 2, 3, and 4 to buttons from left to right. Additionally, a small LED was

programmed to illuminate for the duration of each '4' keypress, visible to the GoPro camera to facilitate post-hoc video syncing. The U-HID allowed for precise mapping of keypresses to specific functions, facilitating accurate data capture of participants' inputs during the keypress sequences.

The **SMALLRIG 5 in Magic Arm with Clamp** is a versatile mounting solution designed to securely attach monitors, projectors, LED lights, and cameras, such as the GoPro, to various surfaces. The flexible 5-inch arm allows for precise angle adjustments, ideal for positioning the GoPro cameras for optimal recording of hand movements during motor sequence performance. The Clamp component is compatible with the T-slot aluminium structure, enabling a stable and customizable mount on the apparatus frame, which allowed easy adjustments during setup and ensured reliable placement throughout data collection.

Lastly, the **GoPro HERO10 Black Camera** is an advanced action camera equipped with a 23.6-megapixel sensor, enabling high-resolution video capture at up to 5.3K at 60 frames per second. The camera's compact design and image quality make it ideal for recording detailed hand movements. For this study, the camera was set to record at a resolution of 1280x720 and a frame rate of 120 fps.

Note: The GoPro HERO10 automatically splits video files once they reach a size of 4GB. As a result, most participant videos from Study 3 were internally split into two separate files. The methods for programmatically reducing the resolution and "stitching" the split videos back together will be covered in a subsequent section, accompanied by the corresponding Python code.

5.2.3 Inclusion criteria

Participants were recruited if they were $>$ 18 years old had normal or corrected-to-normal vision. Exclusion criteria included a history of severe head trauma associated with loss of consciousness, brain surgery, or surgical procedures to the spinal cord. Participants with upper limb motor impairments or any pain or discomfort of the hands that would prevent them from using a computer keyboard were also excluded. Additionally, individuals who were taking or had taken any prescribed medications as part of treatment or research within the last two weeks (excluding

contraceptives) were not eligible to participate. Those suffering from any neurological or psychiatric diseases were excluded to ensure the integrity of the study's focus on motor sequence learning in a neurologically healthy population.

5.2.4 Participants

Sixty naive right-handed healthy participants (51 women; mean ± SD age 20.1 ± 3.27) gave their written informed consent to participate in the study, which was approved by the UCL Research Ethics Committee. Active musicians were excluded from the study (Abraham and Drory, 2014). The sample size was determined a priori via a power analysis of prior skill learning data collected in my NIH research group using a similar motor sequence task. Of the 60 participants recruited, data from 13 participants were excluded due to technical issues: 6 camera failures, 5 online data collection failures, 1 SD card failure, and 1 keyboard failure. This resulted in a final sample size of 47 participants included in the analysis.

5.2.5 Study Design

Participant Assignment and Rotational Structure

Participants in Study 3 were randomly assigned to one of six rotations, each defining a unique order of motor sequence tasks. In this study, I will refer to these unique orders as "rotations" rather than "conditions" to distinguish them from traditional experimental conditions. The purpose of these rotations was threefold: to balance participants' exposure to various motor sequences, observe skill acquisition across different trial structures, and control for order effects in the results. My design ensures that any observed differences in skill learning are not artifacts of sequence order but rather reflect differences in learning dynamics.

Each rotation consists of a distinct ordering of paired sequences, with each pair composed of an initial sequence followed by a paired sequence in which the initial sequence is embedded. Rotations adjust the order of these sequence pairs by rotating them counter-clockwise across the six possible assignments. Note the initial and paired sequence text in red below and it's 'Rotation' over each assignment. Further, note that each Rotation lists the number of participants in (parentheticals).

- **Rotation 1 (6 Participants):**

  (413, 24134) (241, 12413) (143, 31432)

  (342, 23423) (324, 43241) (234, 32342)

- **Rotation 2 (7 Participants):**

  (241, 12413) (143, 31432) (234, 32342)

  (413, 24134) (342, 23423) (324, 43241)

- **Rotation 3 (9 Participants):**

  (143, 31432) (234, 32342) (324, 43241)

  (241, 12413) (413, 24134) (342, 23423)

- **Rotation 4 (8 Participants):**

  (234, 32342) (324, 43241) (342, 23423)

  (143, 31432) (241, 12413) (413, 24134)

- **Rotation 5 (8 Participants):**

  (324, 43241) (342, 23423) (413, 24134)

  (234, 32342) (143, 31432) (241, 12413)

- **Rotation 6 (9 Participants):**

  (342, 23423) (413, 24134) (241, 12413)

  (324, 43241) (234, 32342) (143, 31432)

Trial Structure and Task Progression

Each sequence pair within a rotation was practised across multiple trials in a structured progression. Participants first performed three trials of the initial sequence followed by three trials of the paired sequence, with subsequent increases in trial counts per pair: seven trials for the next initial-paired sequence set, then eleven trials. After these trial blocks, participants took a five-minute rest.

Following the rest, participants continued with the next set of sequences, starting again with three trials per sequence in the pair. This structured approach pro-

vided multiple opportunities to learn each sequence under controlled trial counts while ensuring balanced exposure to both initial and paired sequences.

Table 5.2 illustrates the sequence progression for Rotation 1, showing the number of trials for each initial and paired sequence. In this table, characters "A" and "B" denote the initial and paired sequences, respectively. Superscript numbers indicate the sequential order of the sequences within Rotation 1 (e.g., A^1 , B^1 represent the first sequence pair, followed by A^2 , B^2 , and so forth). The sequences are practised in blocks of increasing trials: 3 trials each for the first pair, 7 trials each for the second pair, and 11 trials each for the third pair.

	3 Trials each	**7 Trials each**	**11 Trials each**
Initial Sequence	**A[1]** 413	**A[2]** 241	**A[3]** 143
Paired Sequence	**B[1]** 24134	**B[2]** 12413	**B[3]** 31432
Initial Sequence	**A[4]** 342	**A[5]** 324	**A[6]** 234
Paired Sequence	**B[4]** 23423	**B[5]** 43241	**B[6]** 32342

Table 5.2: Trial structure for Rotation 1, illustrating the sequence progression through initial and paired sequences with varied trial counts. Cells shaded in red denote trials before a 5-minute rest, while blue cells indicate trials after the rest period.

Cells shaded in red indicate trials conducted before a mandatory 5-minute rest period, while cells shaded in blue represent trials that occur after the rest. This alternating structure allows participants to revisit different motor sequences in varied trial lengths, ensuring a balanced exposure to all sequences while allowing sufficient rest intervals to promote skill consolidation. This rotation design minimizes ordering effects, helping to isolate learning patterns across different sequences.

5.2.6 Analysis

The following section outlines analysis methods unique to Study 3, introduced due to differences in study design and data complexity from Studies 1 and 2. The subsequent sections provide detailed descriptions of the modified skill measure calculation, UMAP with Hamming distance, and recursive HDBSCAN clustering, used to assess motor sequence learning and synergy use in this study.

Skill Measure Calculation

In Study 3, I introduced a modified skill metric, switching from the correct sequences per second (cs/s) used in Studies 1 and 2 to correct keypresses per second (ckp/s). The new measure for motor skill learning focused on individual keypresses rather than complete sequences, accommodating the changing sequence lengths (from 3 to 5 elements) and diverse sequence orders across rotations.

The calculation of ckp/s involved the following steps:

1. Define specific target sequences for each rotation and parse these sequences as integer arrays.

2. Use `patternDetect` function to evaluate each participant's stream of keypresses against the target sequence, returning a score based on the number of matching pairs.

3. Compute mean correct keypresses per second (ckp/s) for each trial and calculate keypress transition speeds to gauge the smoothness of transitions within each sequence.

This code defines the `patternDetect` function, which calculates correct keypresses per second by identifying pairs in the participant's keypress stream that match the target sequence. A 'ghost element' is appended at the end of the target sequence, allowing the function to treat the sequence circularly, ensuring accurate detection of repeated patterns as participants type the target sequence continuously. Interruptions, or instances where pairs do not match the target, are tracked to account for transitions back to the target sequence, providing a more robust measure of participants' adherence to the target pattern over time.

It is important to note here that although participants typed the 3-element sequences continuously during each 10-second trial, I verified that they adhered to the prescribed sequence order (e.g., 4–1–3) rather than adopting rotated alternatives (e.g., 1–3–4). While the *patternDetect* algorithm was designed to be tolerant to sequence rotations, adherence to the correct starting point was independently validated using output from PsyToolkit, which recorded the raw order of keypresses in

each trial. This allowed us to identify the first keypress in every sequence attempt and confirm whether participants initiated the trial with the correct digit. Trials or participants that failed this check were flagged and excluded from further analysis. This verification ensured that motor synergies analysed for transfer were formed in accordance with the intended sequence structure, not alternative permutations.

Transition Speed and Skill Learning Transfer

In Study 3, I examined transition speeds between consecutive keypresses as a measure of motor skill learning transfer. By analysing combined transition speeds, I assessed participants' ability to generalise skill from 3- to 5-element sequences. In this analysis, transition "speed" was operationalised as transition time - specifically, the time in milliseconds between consecutive keypresses. Although the term speed is colloquially associated with higher values indicating faster responses, my measure reflects the inverse: longer durations between keypresses indicate slower transitions. Importantly, I did not compare overall performance between the 3- and 5-element sequences, but rather isolated the same subsequence (e.g., 2-4-1) embedded within both conditions. This allowed us to measure how the execution timing of a learned sequence changed when embedded in a more complex motor context. Thus, an increase in transition time from the 3- to 5-element condition reflects not a general unfamiliarity with the 5-element task, but a measurable slowdown of a previously learned motor sequence when it is nested within a longer, complex sequence. This distinction is critical for interpreting the meaning of the transition time results and a visualisation of the concept is summarised in Figure 5.2.

The calculation of combined transition speeds involved the following steps:

1. **Define transition pairs**: Specific pairs of keypress events, i.e. transition pairs, were identified based on the target sequences.

2. **Calculate time intervals**: Time intervals between consecutive keypresses were computed to determine individual transition speeds.

3. **Identify valid sequence**: Consecutive transitions matching transition pairs were identified and their speeds summed.

4. **Aggregate data**: Combined speeds were aggregated across trials and participants to evaluate performance trends over time.

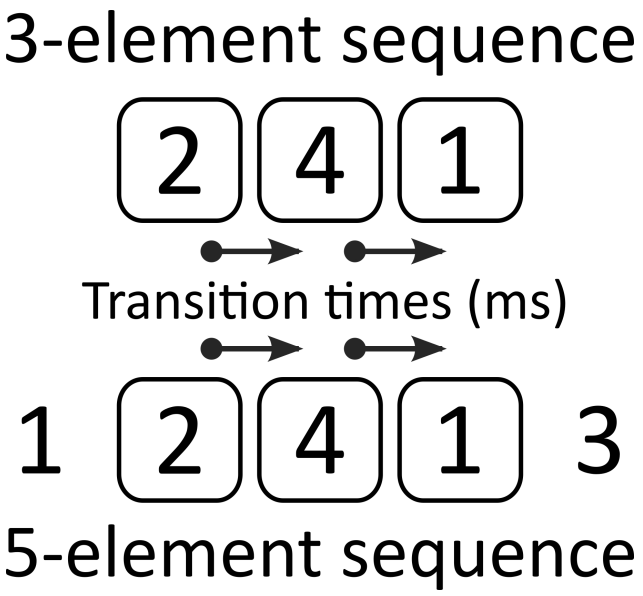

Figure 5.2: Schematic of the transition time comparison used to assess skill transfer. Identical keypress transitions (2→4→1) were isolated from both the final trial of the trained 3-element sequence (2-4-1) and the corresponding subsequence embedded within a paired 5-element sequence (1-2-4-1-3). Arrows denote the inter-key intervals measured in milliseconds. This approach allows evaluation of whether execution speed for a learned subsequence is preserved when recontextualised within a more complex motor task.

This process identifies consecutive keypress events matching transition pairs and calculates their combined transition speed by summing the time intervals of relevant transitions. The resulting list represents the execution speed of sequences. By analysing these combined speeds, I evaluated participants' ability to transfer and generalise motor skills. This approach extends beyond accuracy metrics to emphasise the temporal dynamics of motor learning.

UMAP with Hamming Distance

UMAP with Hamming distance was implemented to explore clustering patterns at the group level, rather than individually as in Studies 1 and 2. This methodological change was necessary due to the large data set, which exceeded the sample counts of Study 1 by over 20-fold, reaching up to 907,200 samples in Rotation 3 alone.

The UMAP embedding with Hamming distance allowed for the incorporation of one-hot encoded features representing trial groupings between initial and paired

sequence trials. This allowed for a synergy analysis between initial and paired trial sets while not overextending the search for transferred synergies beyond each tested set of initial and paired sequences. This process involved:

1. Wavelet-transforming pose data to extract spatial-temporal features, followed by UMAP dimensionality reduction.

2. Using Hamming distance and one-hot encoding to categorize and project pose data into two-dimensional space, allowing loose grouping of specific 3-element and paired 5-element sequences within rotations.

```
mapper = umap.UMAP(n_neighbors=8, n_components=2, min_dist
    =0.3, metric='hamming', init='random')
embed = mapper.fit_transform(proj_with_one_hot)
```

This code sets up the UMAP embedding using Hamming distance, where `proj_with_one_hot` represents the combined wavelet-transformed features and one-hot encoded data across trials in the rotation.

Recursive-HDBSCAN

I used recursive HDBSCAN clustering on the UMAP-embedded data across all participants in a rotation. By recursively applying HDBSCAN on initially discovered clusters, I could identify granular motor synergy patterns that emerge as participants transition from simpler to more complex sequences.

This clustering approach involved:

1. Running initial HDBSCAN clustering on the full UMAP-embedded data set for each rotation.

2. Selecting the original clusters and performing HDBSCAN recursively on each to identify sub-clusters that capture specific synergy patterns.

3. Assigning unique cluster labels based on these refinements, showing detailed and stable clusters that reflect motor synergies across all trials.

```
# Run initial HDBSCAN
clusterobj = hdb_clustering.hdb_scan(embed, 1500, 20,
    selection='leaf', cluster_selection_epsilon=0.05)
labels = clusterobj.labels_
probabilities = clusterobj.probabilities_

# Recursive HDBSCAN on original clusters
for cluster_id in original_clusters:
    selected_cluster_mask = labels == cluster_id
    selected_points = embed[selected_cluster_mask]
    refined_clusterobj = hdb_clustering.hdb_scan(
        selected_points, max(len(selected_points) // 35, 10),
        10, selection='leaf', cluster_selection_epsilon=0.01)
```

Recursive-HDBSCAN procedure allows for detailed exploration of motor synergies within each rotation by iteratively refining clusters based on group-level data.An example of the recursive-HDBSCAN output, as well as a visualisation of how Hamming distance affects the spatial arrangement of the UMAP embedded space can be found in Figure 5.2.

In the next section, I present the results of Study 3, with one section for skill performance (e.g. correct keypresses per second and transition times) and permutation tests across all rotations, showing how motor synergies do or do not transfer across trials and motor sequences. Following this, I detail population-level results - in terms of the kinematic differences between 3- and 5-element motor sequences - across all rotations to describe motoric trends in motor sequence learning transfer. Finally, I provide a summary of results in Study 3 to consolidate the key findings.

5.3 Results

5.3.1 Skill Performance and Learning

Performance, as indexed by correct keypresses per second (ckp/s), improved consistently across sequences and rotations, with some sequences reaching a plateau

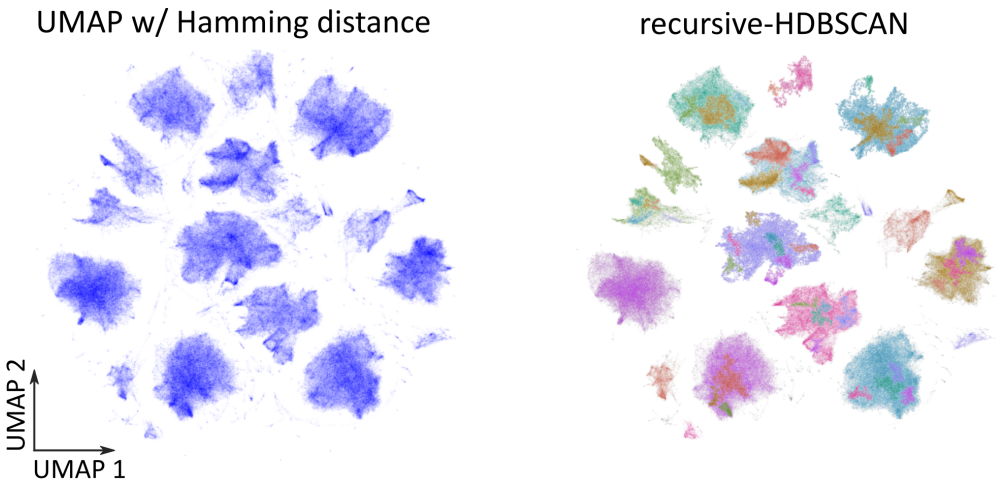

Figure 5.3: (Left) UMAP with Hamming distance shows the two-dimensional projection of wavelet-transformed pose data from all trials in Rotation 6, where each point represents a pose in one video frame. (Right) Recursive-HDBSCAN clustering identifies discrete motor synergies as coloured clusters and sub-clusters. The full pipeline comprises three stages: (1) spatial–temporal embedding of wavelet transformed data, (2) UMAP dimensionality reduction using Hamming distance on one-hot encoded trial groupings, and (3) recursive application of HDBSCAN to extract granular, synergy behaviour clusters.

after 7 or 11 trials. However, closer examination of transition times between shared 3-element sequences in simple and complex tasks revealed a distinct slowing, indicating altered motor execution in the complex task context. Results from each Rotation are available in Figure 5.3 and Figure 5.4.

It is important to emphasise: To first assess how practised transitions generalised within more complex contexts, I compared the average keypress transition times for identical transition pairs (e.g., 2→4, 4→1, 1→3) in both 3-element sequences and their embedded locations in 5-element sequences. The resulting deltas reflect the *increase* in transition time when these same transitions were executed in the more complex context. These slower transition times indicate that even well-practised transitions became more effortful when embedded in novel or longer sequences, suggesting that synergy transfer was incomplete.

In **Rotation 1**, ckp/s improved across most sequences, such as 24134 and 31432, which showed increases of 1.33 ± 0.39 and 1.00 ± 0.37, respectively. Transition time deltas revealed slower execution of known transitions in the 5-element sequence: for instance, 241 to 12413 showed a delta of 267.91 ms ($p = 0.0003$), indi-

cating degraded fluency. **Rotation 2** showed similar learning gains, with sequences 12413 and 342 increasing by 1.07 ± 0.47 and 1.25 ± 0.36. Again, transitions such as 241 to 12413 were significantly slower by 453.73 ms (p = 0.0014), highlighting reduced transition efficiency when context complexity increased. In **Rotation 3**, skill improvements were seen in sequences 23423 and 43241, with differences of 1.42 ± 0.42 and 1.12 ± 0.43. Sequence 241 increased by 1.30 ± 0.48. The transition from 324 to 43241 had a delta of 302.77 ms (p = 0.0217), while 241 to 12413 had a non-significant delta of 528.60 ms (p = 0.1212), suggesting mixed transfer effects.

Rotation 4 featured strong ckp/s gains in sequences like 43241 (1.50 ± 0.21) and 234 (1.24 ± 0.33). The transition time delta from 342 to 23423 was 257.84 ms (p = 0.0011), while 241 to 12413 showed a delta of 416.79 ms (p = 0.0009), again pointing to slowed transitions despite prior practice. In **Rotation 5**, sequences 143 and 324 improved by 1.38 ± 0.31 and 1.36 ± 0.43. Sequence 413 showed the greatest increase (1.74 ± 0.30). The transition from 234 to 32342 was slower by 370.98 ms (p = 0.0003), and from 241 to 12413 by 319.19 ms (p = 0.0006). Finally, **Rotation 6** participants demonstrated increases in 43241 (1.52 ± 0.38) and 31432 (1.23 ± 0.37), with 24134 showing the greatest improvement (1.58 ± 0.43). Transition time deltas were again substantial: 234 to 32342 was slower by 390.21 ms (p = 0.0001), and 241 to 12413 by 379.17 ms (p = 0.0042).

5.3.2 Permutation Tests

Permutation tests were conducted to assess whether motor synergies developed during the final trial of the 3-element sequence transferred to the first trial of the paired 5-element sequence. For each of the 36 sequence comparisons across six rotations, 100,000 permutations were performed to generate empirical null distributions of JSD values. The overwhelming majority of comparisons yielded significant divergence, indicating non-transfer of motor synergies across sequence contexts.

In **Rotation 1**, strong divergence was observed in all sequences. Notable results included 324 to 43241 (JSD = 0.6897, p = 1.0e-04), 234 to 32342 (JSD = 0.6857, p = 1.2e-04), and 241 to 12413 (JSD = 0.6830, p = 2.3e-04). The least significant result, 342 to 23423, still reached significance (JSD = 0.6349, p = 1.291e-

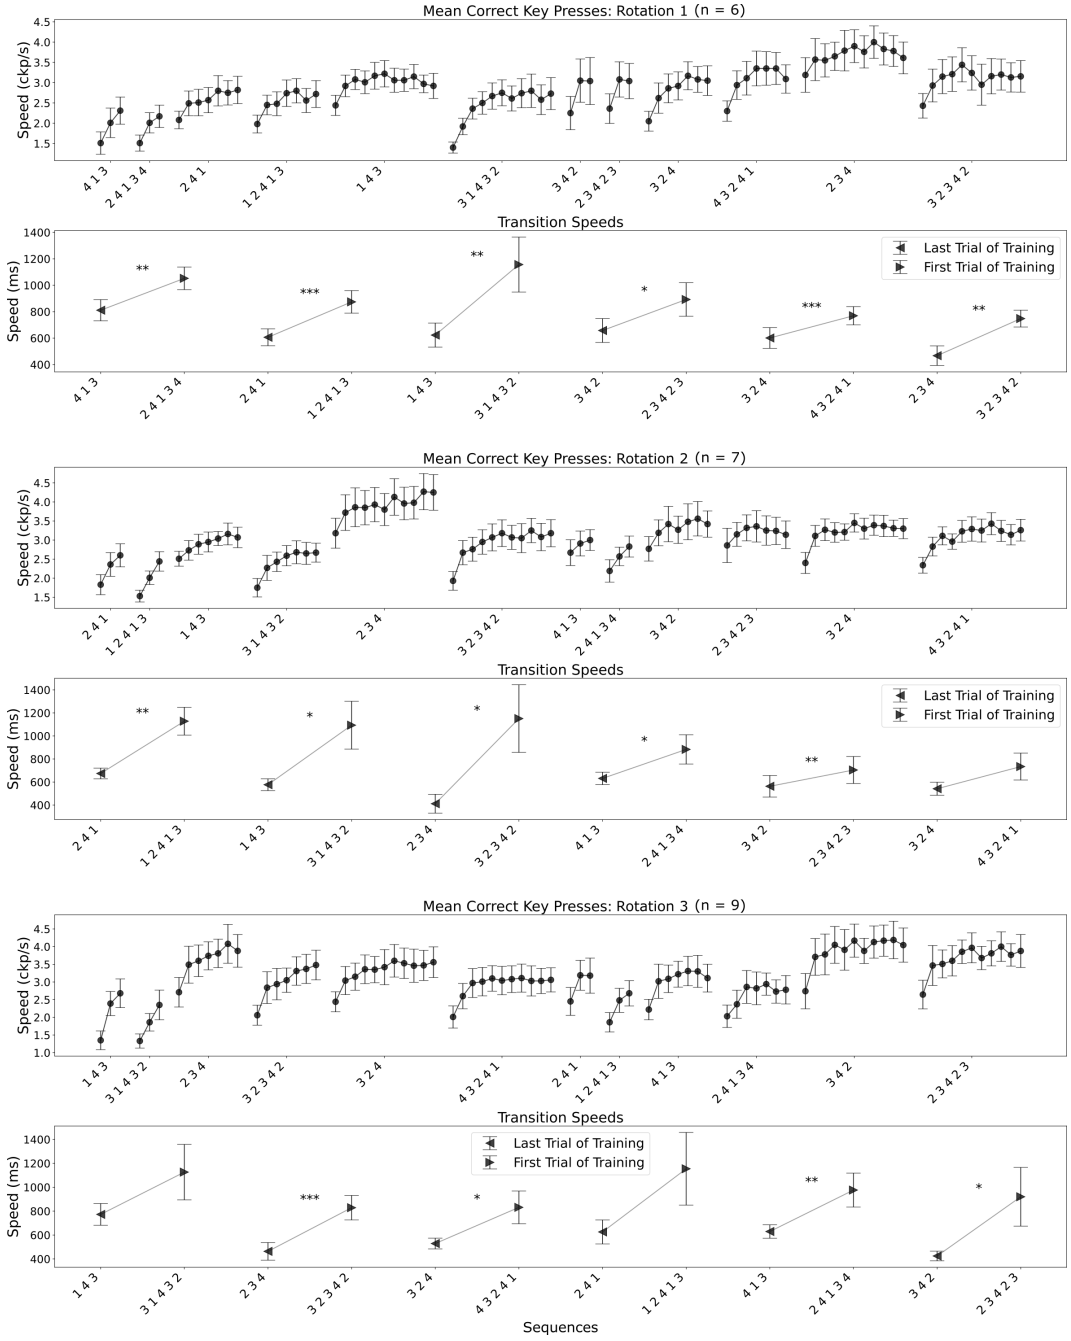

Figure 5.4: Skill performance and transition speed comparisons for Rotations 1–3. Top panels depict mean ckp/s across all trials for each sequence, with error bars indicating ±SEM. Each cluster of points corresponds to a different sequence, revealing improvements in ckp/s as participants progressed. Bottom panels show average keypress transition speeds (ms) between identical 3-element subsequences embedded in both 3- and paired 5-element sequences. Specifically, transition times from the last trial of the 3-element sequence are compared to the corresponding transitions within the first trial of the 5-element sequence.

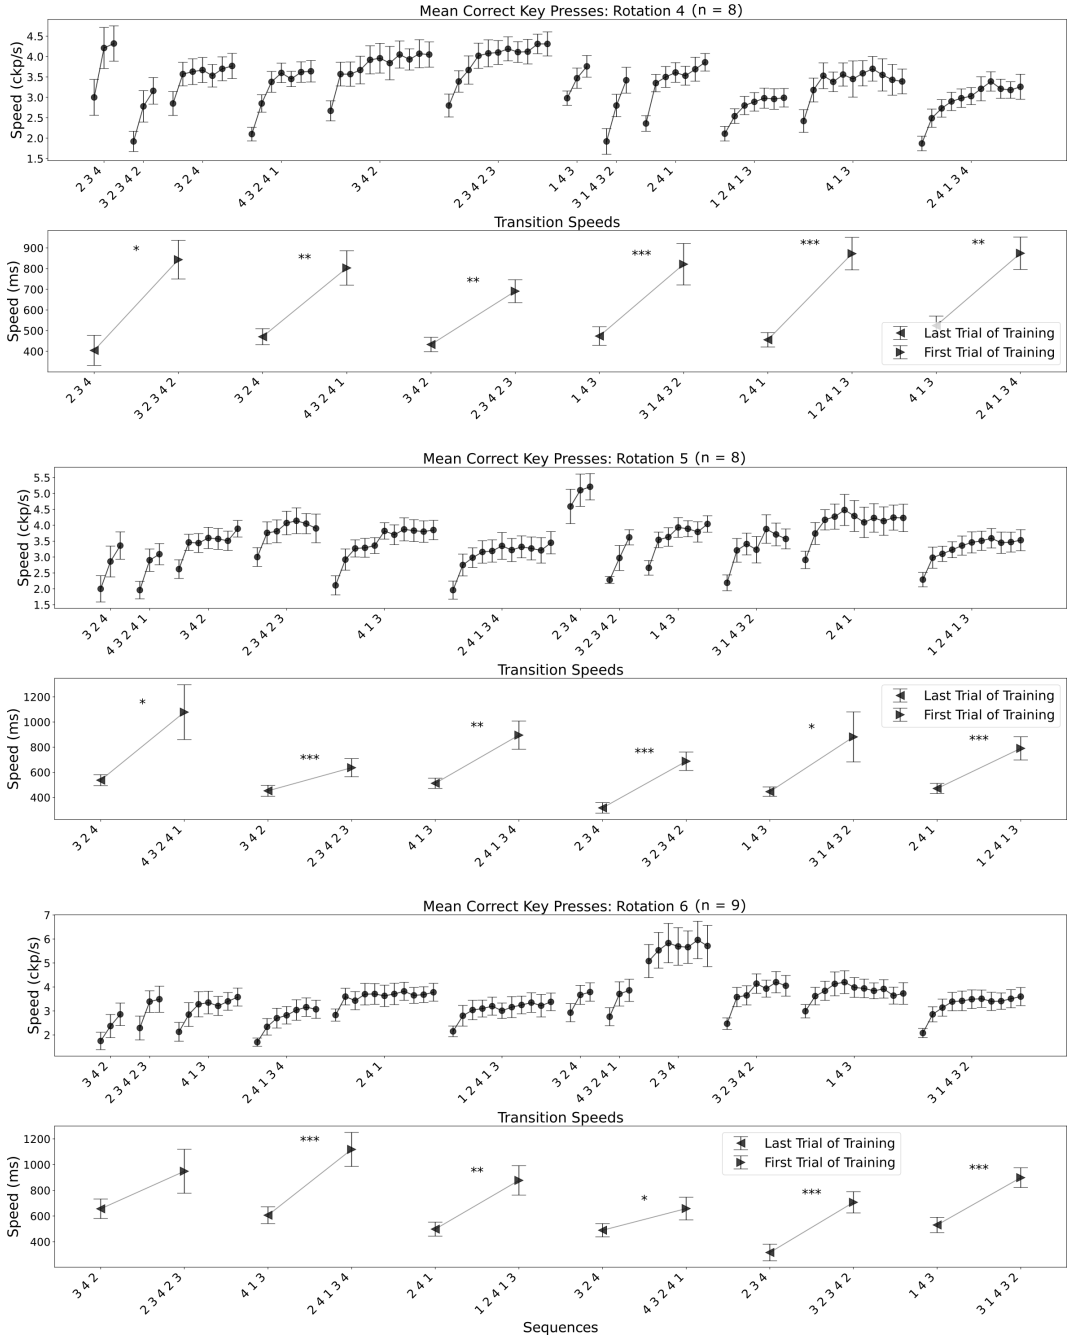

Figure 5.5: Skill performance and transition speed comparisons for Rotations 4–6. Top panels depict mean ckp/s across all trials for each sequence, with error bars indicating ±SEM. Each cluster of points corresponds to a different sequence, revealing improvements in ckp/s as participants progressed. Bottom panels show average keypress transition speeds (ms) between identical 3-element subsequences embedded in both 3- and paired 5-element sequences. Specifically, transition times from the last trial of the 3-element sequence are compared to the corresponding transitions within the first trial of the 5-element sequence.

02). **Rotation 2** showed comparable findings. Sequences 143 to 31432 (JSD = 0.6851, p = 1.61e-03), 413 to 24134 (JSD = 0.6825, p = 3.05e-03), and 234 to 32342 (JSD = 0.6854, p = 3.95e-03) all demonstrated robust differences. The weakest, 342 to 23423, again remained significant (JSD = 0.6873, p = 1.049e-02).

In **Rotation 3**, sequence 241 to 12413 (JSD = 0.6844, p = 1.0e-04) and 234 to 32342 (JSD = 0.6869, p = 5.5e-04) showed strong divergence. Sequence 413 to 24134 reached only marginal significance (JSD = 0.6834, p = 3.038e-02), but remained consistent with the broader trend. **Rotation 4** continued the pattern, with 234 to 32342 (JSD = 0.6763, p = 9.2e-04), 324 to 43241 (JSD = 0.6842, p = 3.3e-03), and 143 to 31432 (JSD = 0.6845, p = 1.72e-03) all showing significant divergence. The weakest results again came from 241 to 12413 (JSD = 0.6846, p = 1.388e-02) and 413 to 24134 (JSD = 0.6846, p = 1.844e-02).

In **Rotation 5**, the sequences 324 to 43241 (JSD = 0.6874, p = 1.0e-04), 234 to 32342 (JSD = 0.6876, p = 6.0e-05), and 342 to 23423 (JSD = 0.6829, p = 9.7e-05) showed the strongest divergence. Other sequences such as 413 to 24134 (JSD = 0.6847, p = 4.7e-03) remained significant. Sequence 241 to 12413 (JSD = 0.6793, p = 1.903e-02) was again among the weaker - but still significant - results. **Rotation 6** replicated the pattern. Sequences 342 to 23423 (JSD = 0.6876, p = 2.3e-04), 241 to 12413 (JSD = 0.6838, p = 6.3e-04), and 413 to 24134 (JSD = 0.6854, p = 5.5e-04) all demonstrated substantial divergence. The least significant result was 234 to 32342 (JSD = 0.6833, p = 1.082e-02).

Study 3 results confirm a consistent pattern of non-transfer in motor synergies across all Rotations. Despite variability in sequence content, participant group, and learning duration, the final trial of the 3-element sequence consistently differed from the initial trial of its paired 5-element sequence. This supports the central finding that synergies formed in one task do not immediately generalise to a more complex task, even when sequence components are embedded.

5.3.3 Population Results

Figures 5.8, 5.9, and 5.10 show results regarding the gain or loss of velocity, acceleration, or overlap when transitioning from 3-element to 5-element sequences. In

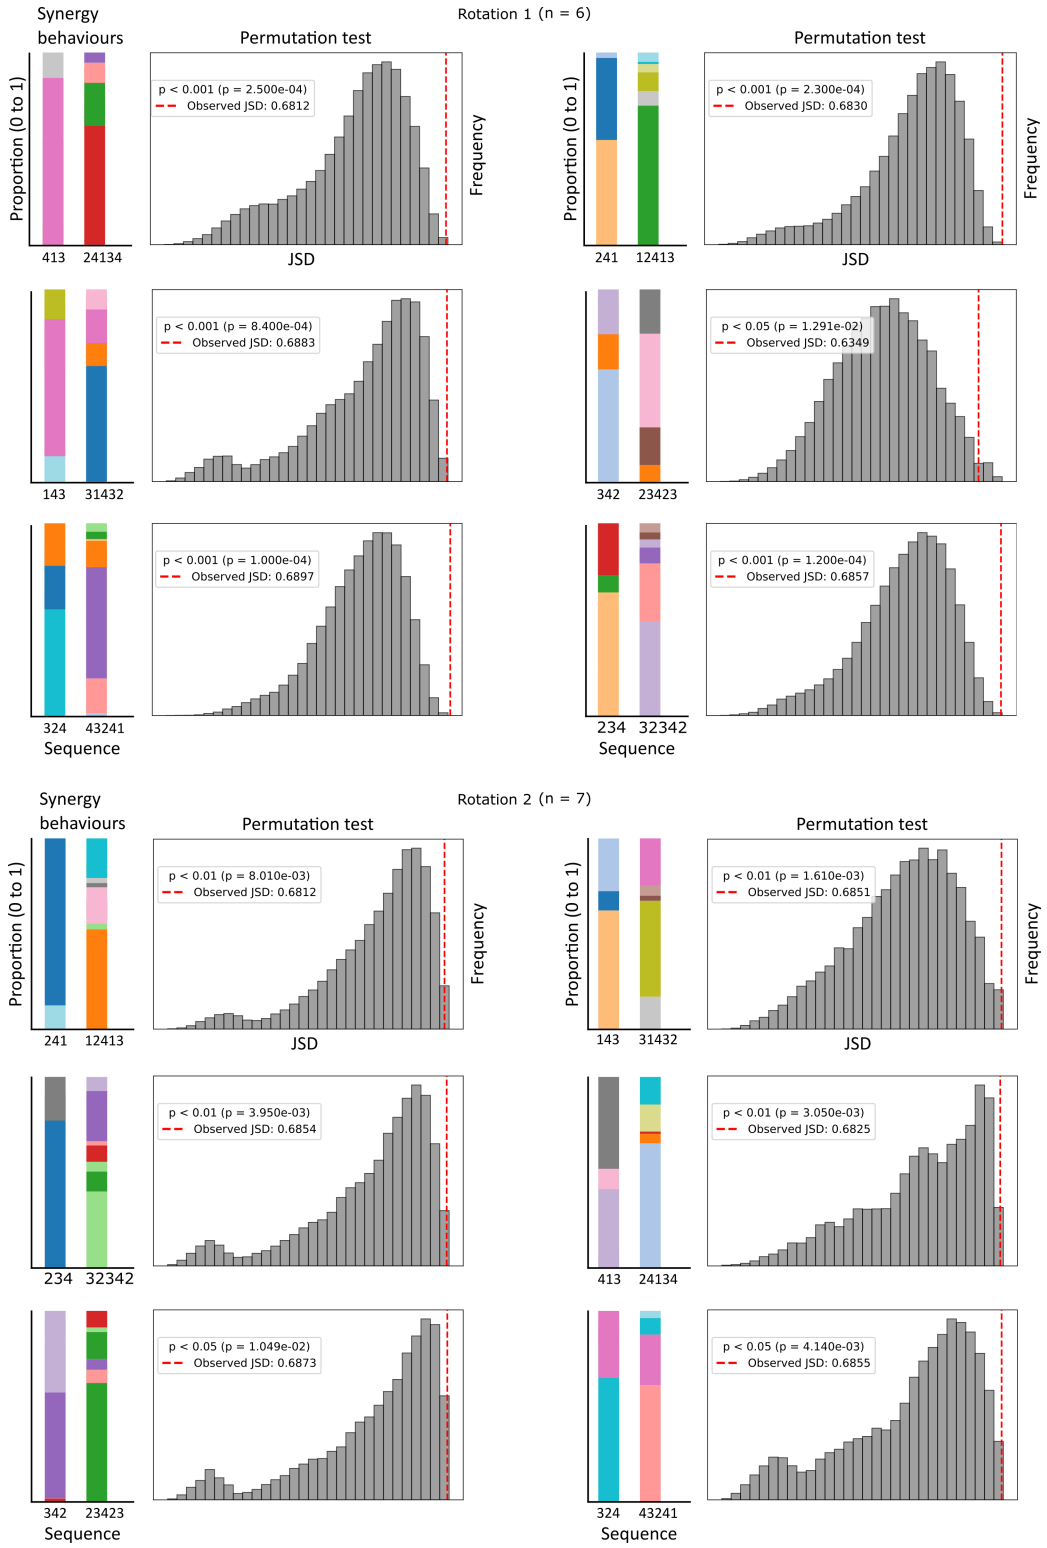

Figure 5.6: Any two coloured bars show synergy distributions from the final trial of a 3-element sequence and the first trial of its corresponding 5-element sequence. Colours indicate distinct synergy clusters. Adjacent histograms display null distributions of JSD values from permutations; red dashed lines mark the observed JSD, with p-values indicating the proportion of permutations yielding a JSD of equal or greater magnitude.

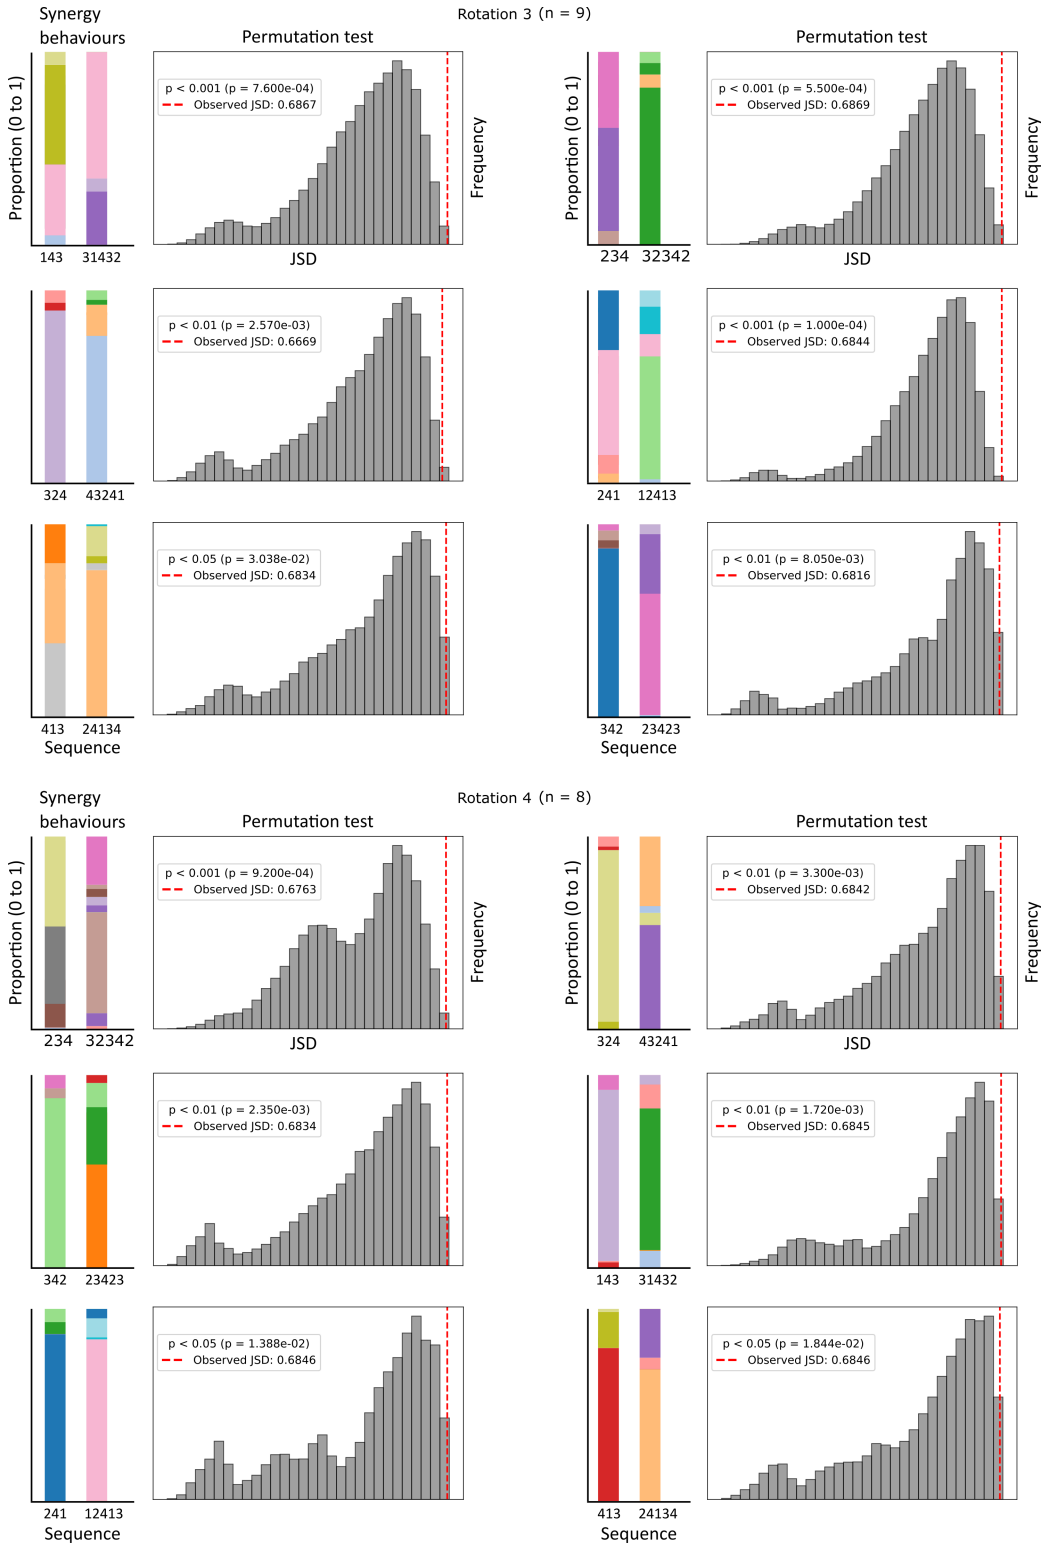

Figure 5.7: Any two coloured bars show synergy distributions from the final trial of a 3-element sequence and the first trial of its corresponding 5-element sequence. Colours indicate distinct synergy clusters. Adjacent histograms display null distributions of JSD values from permutations; red dashed lines mark the observed JSD, with p-values indicating the proportion of permutations yielding a JSD of equal or greater magnitude.

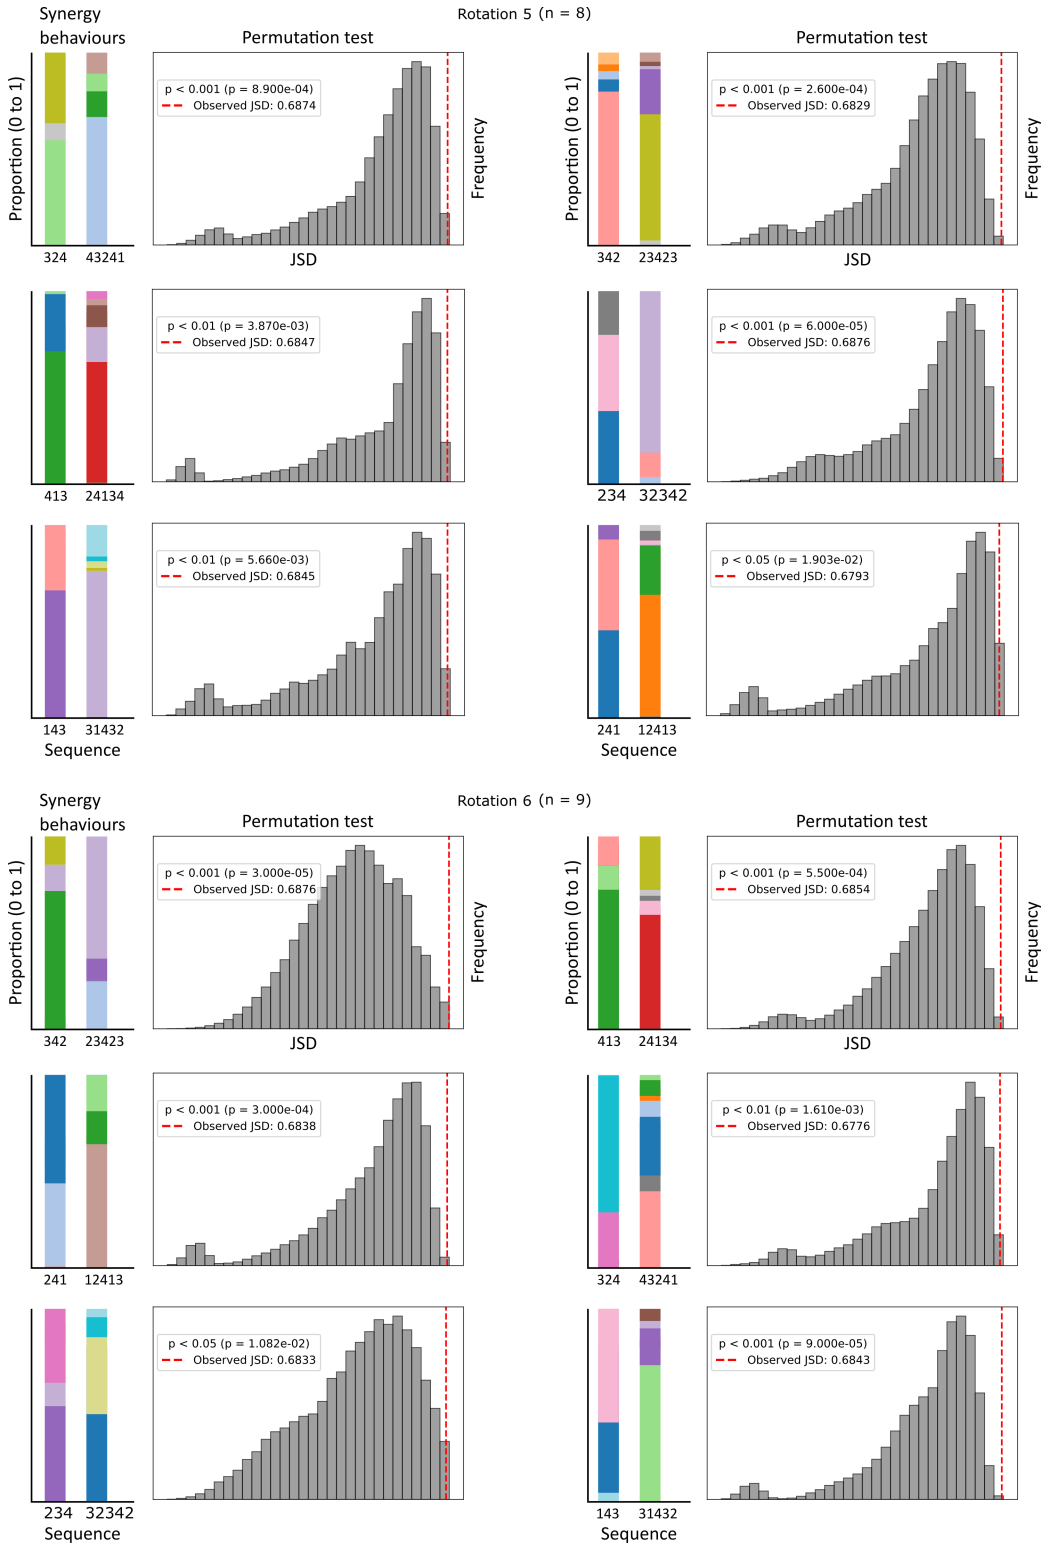

Figure 5.8: Any two coloured bars show synergy distributions from the final trial of a 3-element sequence and the first trial of its corresponding 5-element sequence. Colours indicate distinct synergy clusters. Adjacent histograms display null distributions of JSD values from permutations; red dashed lines mark the observed JSD, with p-values indicating the proportion of permutations yielding a JSD of equal or greater magnitude.

each figure, each plot has a title-detail written at the top of the plot, e.g. ”Sequence: 342”, which is in reference to the 3-element sequence that was the initial sequence and then embedded into the paired, 5-element sequence. In terms of velocity, there were significant differences between any last trial of a 3-element sequence and the first trial of the paired 5-element sequence across participants. For instance, Sequence 342 showed a marked difference in velocity with a higher mean for the last trial of the 3-element sequence (2.181 ± 0.246) compared to the first trial of the 5-element sequence (1.700 ± 0.224), resulting in a large effect size (Cohen's d = 0.728, p = 8.99e-06). The same trend was observed for Sequences 413 (Cohen's d = 0.782, p = 2.62e-06) and 241 (Cohen's d = 0.799, p = 1.73e-06), all of which showed significant velocity advantages in the last trial of the 3-element sequence. The greatest difference was seen in Sequence 234, where the last trial of the 3-element sequence (2.460 ± 0.234) greatly exceeded the first trial of the 5-element sequence (1.396 ± 0.177), resulting in the highest Cohen's d of 1.042 (p = 5.51e-09). Overall, the last trial of the 3-element sequence consistently demonstrated higher velocities across sequences (Overall d = 1.350, p = 4.45e-12).

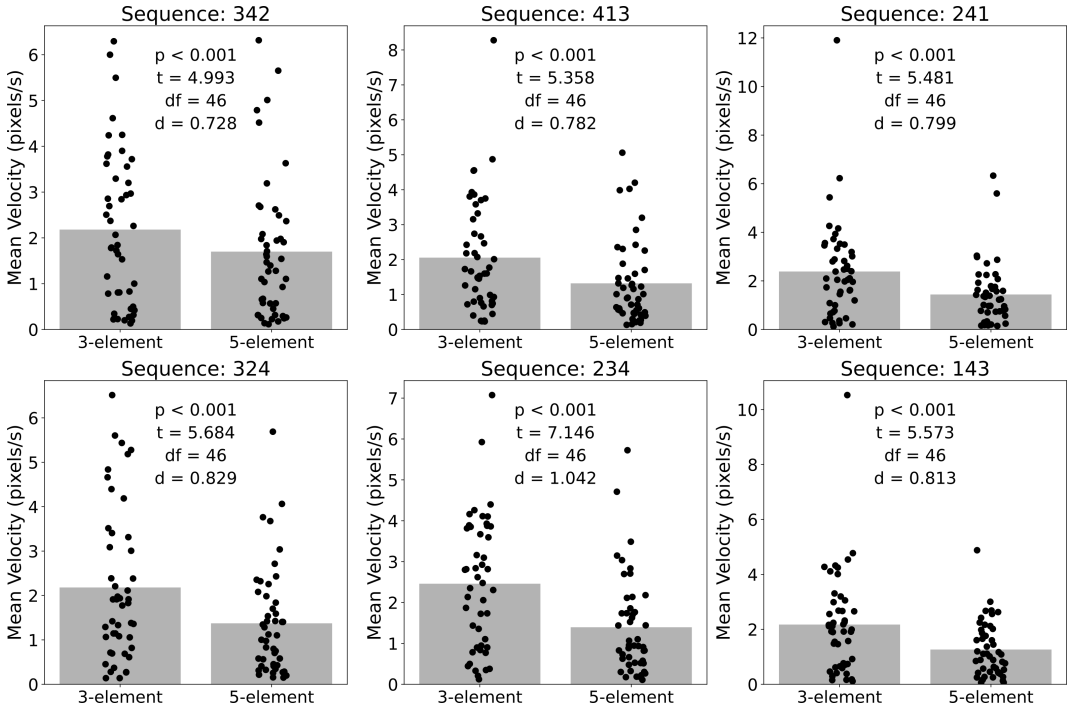

Figure 5.9: Changes In Velocity Between Sequences.

Acceleration data presented less striking differences, with some sequences failing to reach significance. For instance, in Sequence 342, while the acceleration in the last trial of the 3-element sequence (0.718 ± 0.109) was higher than in the first trial of the 5-element sequence (0.592 ± 0.118), the difference was not statistically significant (p = 2.12e-01, Cohen's d = 0.184). Similarly, Sequence 241 showed no significant difference in acceleration (p = 5.32e-01, Cohen's d = 0.092). However, some sequences did show notable differences. Sequence 324 had a significant difference, with a mean acceleration in the last trial of the 3-element sequence of 0.878 ± 0.150 and 0.599 ± 0.104 in the first trial of the 5-element sequence (p = 1.60e-02, Cohen's d = 0.365), and Sequence 143 also showed a meaningful difference with a Cohen's d of 0.434 (p = 4.63e-03). When averaged across all sequences, the last trial of the 3-element sequence had a modest advantage in acceleration (Overall d = 0.467, p = 2.49e-03), though the effect sizes were small to medium.

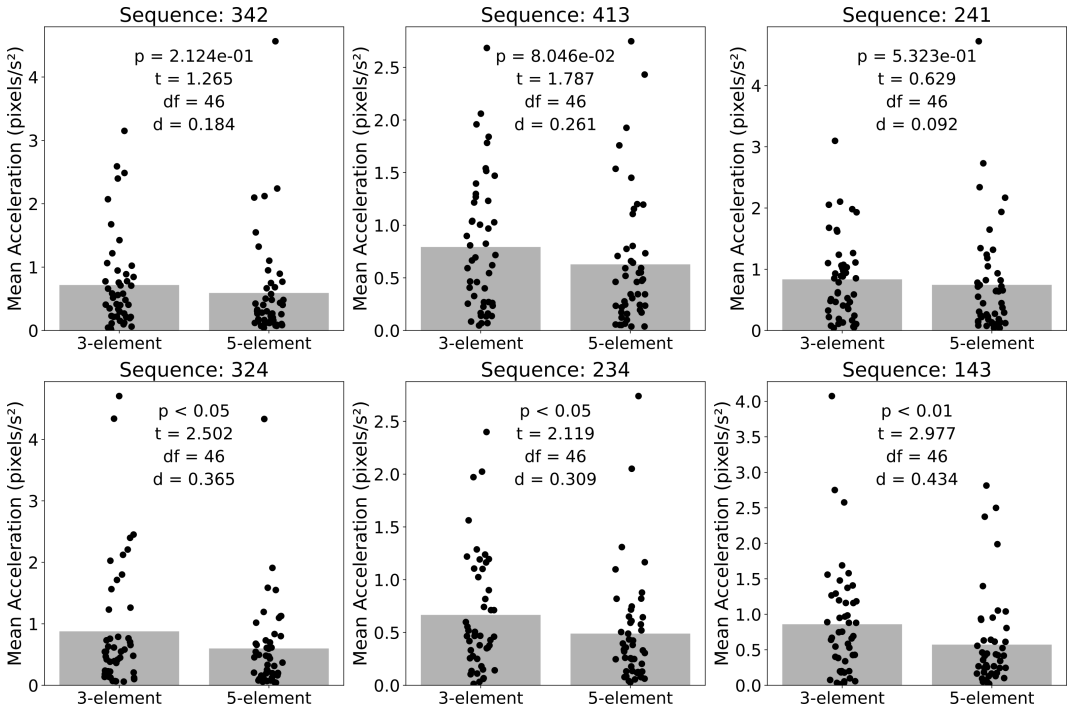

Figure 5.10: Changes In Acceleration Between Sequences.

Overlap data exhibited the most variability across sequences. In some cases, such as Sequence 342, there was no significant difference in overlap between the last trial of the 3-element sequence and the first trial of the 5-element sequence (p

= 5.78e-01, Cohen's d = -0.082). In contrast, certain sequences revealed substantial differences. For instance, in Sequence 234, the overlap in the last trial of the 3-element sequence (19.489 ± 4.012) was significantly higher than in the first trial of the 5-element sequence (2.128 ± 0.570), leading to a large effect size (Cohen's d = 0.658, p = 4.43e-05). Sequence 413 also demonstrated a marked difference, with a Cohen's d of 0.517 (p = 9.09e-04). Overall, the last trial of the 3-element sequence consistently showed greater overlap (Overall d = 0.804, p = 1.56e-06).

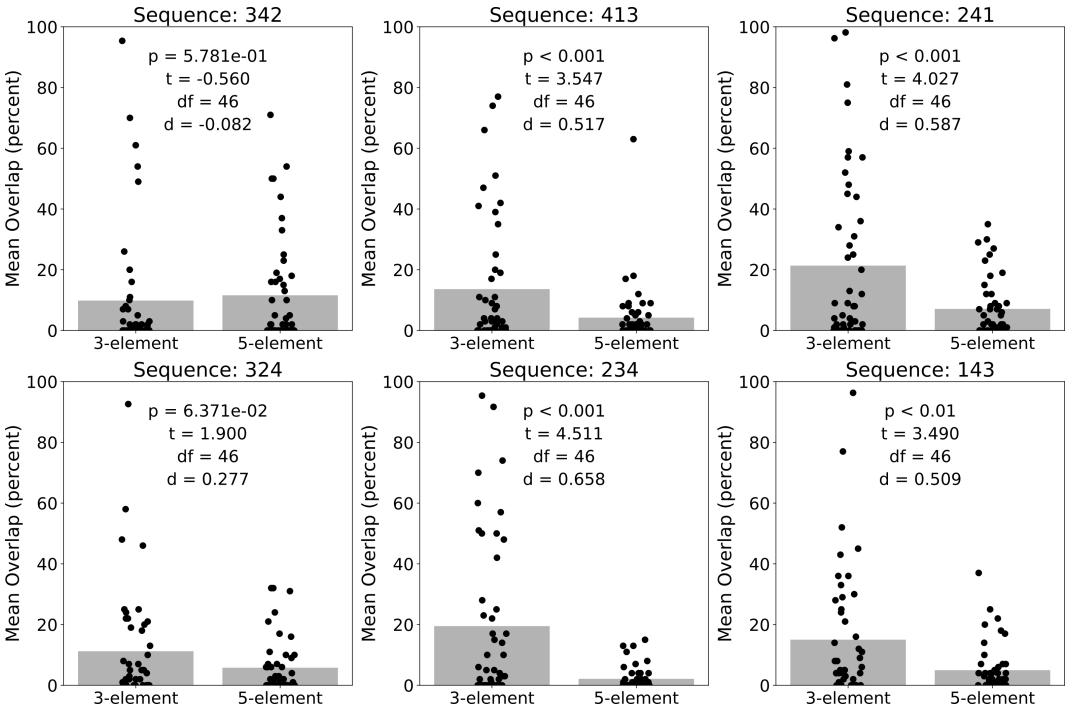

Figure 5.11: Changes In Overlap Between Sequences.

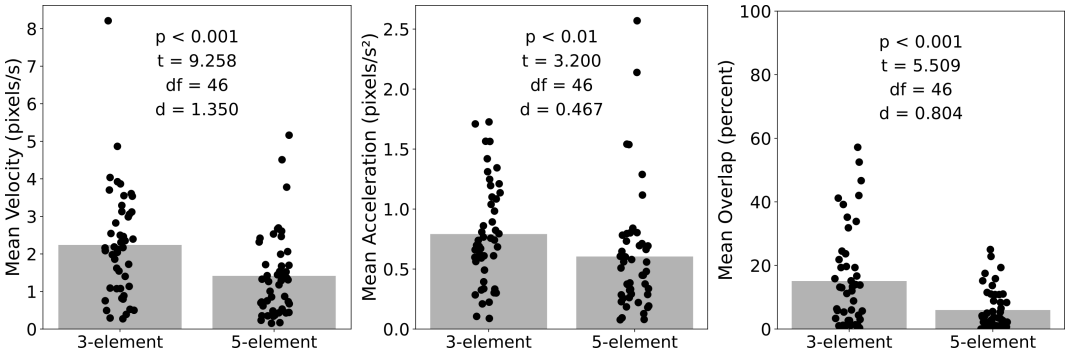

Figure 5.12: Summary of Changes Between Sequences.

If all data from each sequence are collapsed into one summary plot for each

metric – velocity, acceleration, and overlap – the population differences are clear. In Figure 5.11 (Left), there is a significant difference between the 3-element and 5-element sequence velocities, with a t-value of 9.258, degrees of freedom (df) = 46, and Cohen's d = 1.350, indicating a strong effect size ($p < 0.001$). The mean velocity was notably higher in the 3-element sequences, suggesting participants performed faster in these shorter, simpler sequences than in the more complex 5-element sequences. In terms of mean acceleration (Middle), a statistically significant difference was also detected, though with a smaller effect size (Cohen's d = 0.467) and a t-value of 3.200 ($p < 0.01$, df = 46). The 3-element sequences exhibited slightly higher mean acceleration than the 5-element sequences, implying that participants adjusted their movement initiation per keypress in response to the increased complexity in the 5-element sequence. Lastly, mean overlap (Right) showed substantial variation, with a marked difference between the two sequence element lengths (t = 5.509, Cohen's d = 0.804, $p < 0.001$, df = 46). The higher overlap values in the 3-element sequence suggest that participants exhibited greater temporal overlap in their actions when executing the simpler sequence. This decrease in overlap with the introduction of the 5-element sequence indicates a shift towards more segmented movements, potentially reflecting the increased cognitive and motor demands associated with the added sequence complexity.

5.3.4 Disagreement Between Transition Speeds and Kinematics

This section presents example time series of motor sequence performance in cases where transition speeds between the final trial of a 3-element sequence and the first trial of its paired 5-element sequence showed no significant differences, yet observed JSD values between motor synergy distributions indicated significant differences. These examples highlight instances where the timing between keypresses might suggest successful skill transfer, but the pose time series and JSD analyses show notable differences, suggesting that the underlying motor synergies did not transfer between the sequences. By examining each pair of sequences within the shaded regions of the time series, discrepancies between skill learning and execution can be observed.

Rotation 3: Sequence 143 and 31432

In Figure 5.22, I compare the final trial of sequence 143 (top) and the first trial of sequence 31432 (bottom) in Rotation 3. Transition speed analysis between trials showed no significant difference (mean speeds: 772.70 ± 91.01 ms vs. 1126.36 ± 232.26 ms; t = -2.02, p = 0.0741, d = -0.62), which could suggest a transfer in skill learning. However, JSD analysis revealed a significant difference (observed JSD = 0.6867, p = 7.6e-04), indicating distinct motor synergies between these sequences.

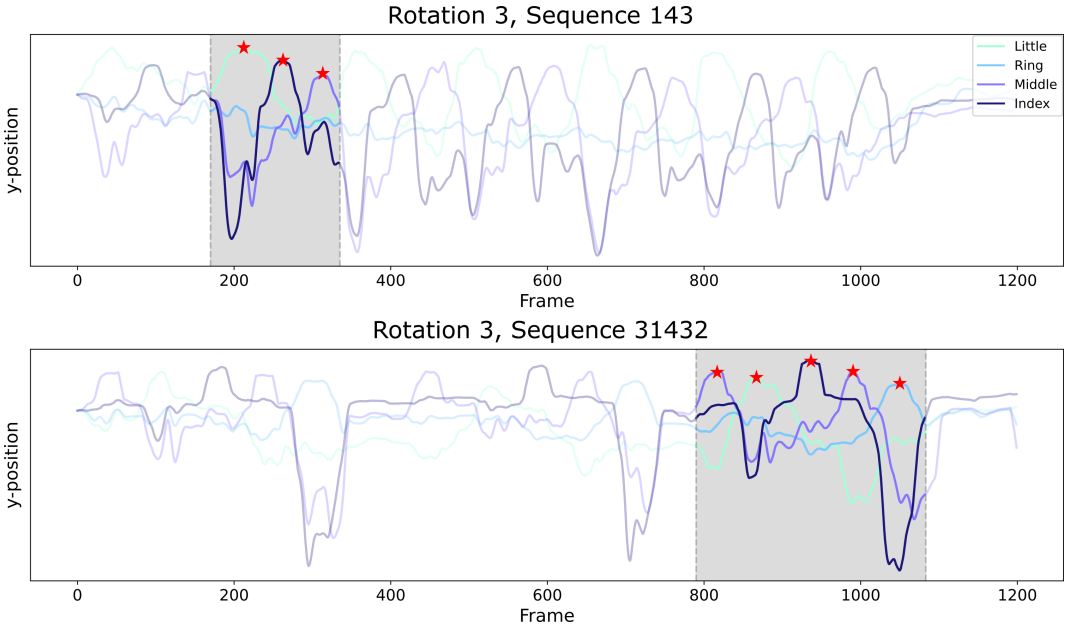

Figure 5.13: Pose Differences Between Sequences, Examples A & B.

Within the shaded region, pose data reveal subtle yet important shifts. In sequence 143, the index finger shows a stable downward movement for each keypress, while the ring finger's motion remains relatively consistent. In contrast, sequence 31432 shows broader, less uniform movements, particularly in the index and ring fingers, suggesting a reconfiguration of motor patterns for the 5-element sequence. This discrepancy between timing and kinematic patterns emphasizes how the JSD metric captures motor adjustments that transition speeds alone may overlook.

Rotation 3: Sequence 241 and 12413

Figure 5.23 compares sequence 241 (top) with sequence 12413 (bottom) in Rotation 3. Transition speeds between these trials were also non-significant (mean speeds:

625.91 ± 100.47 ms vs. 1154.51 ± 304.05 ms; t = -1.73, p = 0.1212, d = -0.74), but JSD analysis showed a significant difference (observed JSD = 0.6844, p = 1.0e-04).

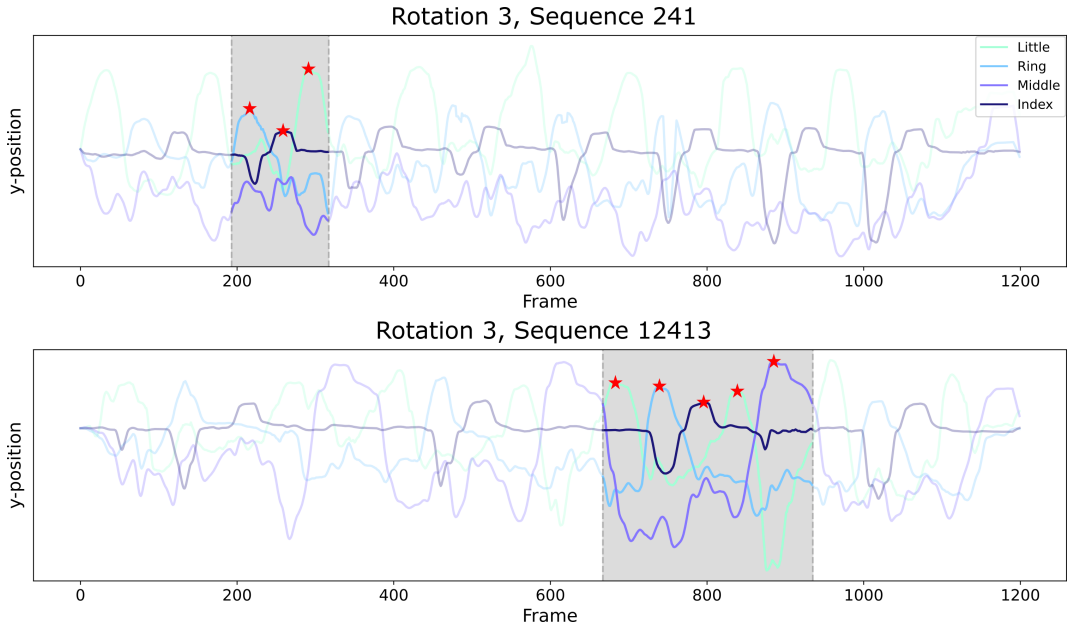

Figure 5.14: Pose Differences Between Sequences, Examples C & D.

The shaded area reveals different kinematic profiles between sequences. In sequence 241, the index and ring fingers move in steady, consistent trajectories with clear downward peaks on each keypress. In sequence 12413, however, the peaks become less consistent, with larger fluctuations, particularly in the index finger. This variation points to a shift in motor synergies despite similar transition speeds, illustrating the value of JSD in highlighting changes in motor control that timing metrics alone might miss.

Rotation 2: Sequence 324 and 43241

Figure 5.24 contrasts sequence 324 (top) with sequence 43241 (bottom) in Rotation 2. Transition speed analysis indicated no significant difference (mean speeds: 541.92 ± 56.75 ms vs. 734.41 ± 116.80 ms; t = -2.21, p = 0.0548, d = -0.64), yet JSD values were significantly different (observed JSD = 0.6855, p = 4.14e-03).

Within the shaded segment, differences in the kinematic patterns of the middle and ring fingers are shown. In sequence 324, these fingers display smooth, uniform trajectories, while in sequence 43241, their movements are more varied, with

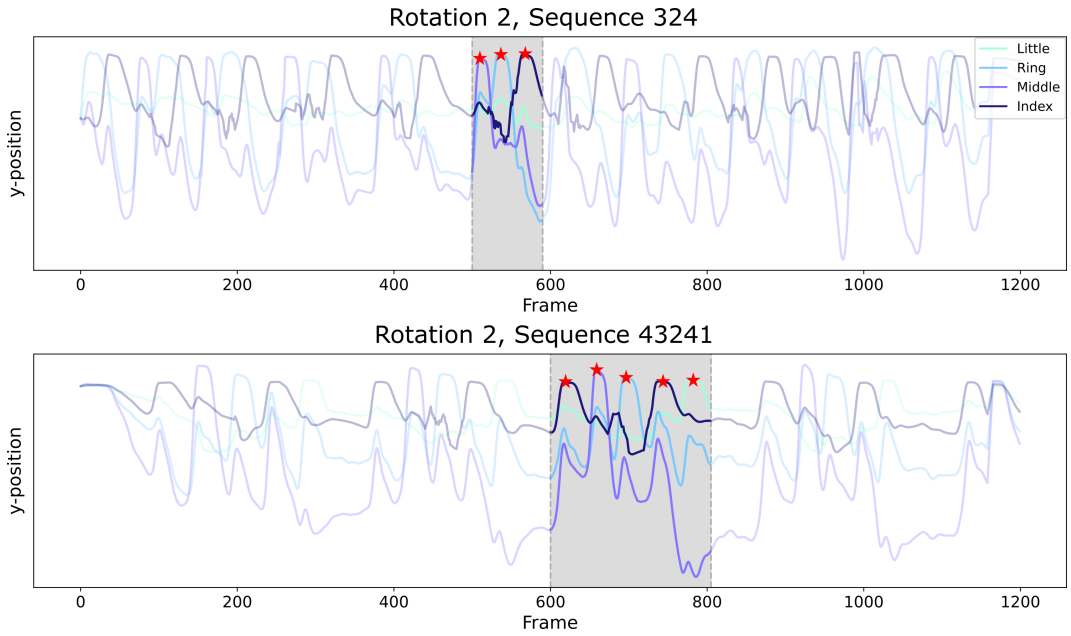

Figure 5.15: Pose Differences Between Sequences, Examples E & F.

pronounced fluctuations and a less stable rhythm. This kinematic shift indicates a change in motor synergy required for the more complex sequence, demonstrating how motor synergy adjustments may remain hidden in transition speed metrics.

Rotation 6: Sequence 342 and 23423

Figure 5.25 compares sequences 342 (top) and 23423 (bottom) in Rotation 6. Transition speeds again showed no significant difference (mean speeds: 656.38 ± 75.80 ms vs. 948.31 ± 170.91 ms; t = -2.19, p = 0.0566, d = -0.67), but JSD analysis indicated a significant difference (observed JSD = 0.6876, p = 3.0e-05).

In the shaded region, the kinematic profiles of the index and middle fingers show distinct patterns between the sequences. Sequence 342 features consistent, sharp downward movements for each keypress, whereas sequence 23423 exhibits larger and more irregular peaks. This change in motor synergy highlights an adaptation to the embedded sequence within the 5-element task, reinforcing the utility of JSD in identifying subtle yet meaningful motor adjustments.

These findings illustrate that even in the absence of significant differences in transition speed, changes in motor synergy distributions and kinematic time series can reveal substantial reorganisation of motor control. Although only 4 out of 36

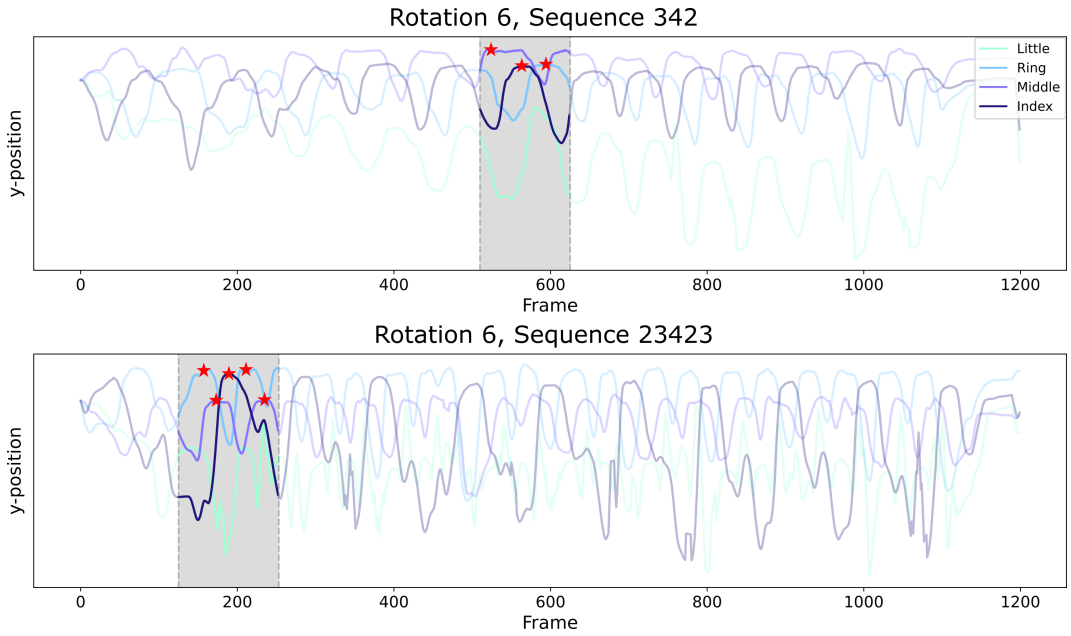

Figure 5.16: Pose Differences Between Sequences, Examples G & H.

sequence pairs showed this dissociation, their presence underscores the limitations of speed-based metrics in capturing transfer. In applied contexts such as rehabilitation or sports, such discrepancies may signal compensatory strategies or incomplete generalisation masked by fluent execution. This highlights the value of incorporating synergy-level analyses alongside conventional performance metrics - a point expanded upon in the Discussion.

5.3.5 Results summary

This section consolidates the findings from Study 3, focusing on rotation-specific results, population-level analyses, and cases where kinematic and transition speed measures disagreed, ultimately emphasizing the lack of motor synergy transfer between subsequent motor sequence skills.

The rotation-specific analyses revealed consistent improvements in participants' performance across rotations. Correct keypresses per second (ckp/s) increased from the first to last trials in each rotation, indicating skill acquisition over time. Transition speeds also showed significant reductions as participants progressed through the trials, suggesting enhanced fluency in executing keypress sequences. Despite these improvements in timing, permutation tests revealed sig-

nificant differences in motor synergy distributions between the final trial of each 3-element sequence and the first trial of the paired 5-element sequence. These observed JSD values indicated substantial divergence in motor synergy patterns between sequence types, supporting the notion that motor synergies developed in one skill do not transfer directly to a subsequent, more complex skill.

Although the primary hypothesis was that synergy transfer would be more likely after brief practice (3 trials) than after longer practice (7 or 11 trials), this pattern was not statistically supported. Permutation tests on trial-level JSD values revealed no significant transfer of motor synergy from the final 3-element trial to the first 5-element trial in any condition. While this null result technically aligns with the hypothesis for 7- and 11-trials, it also contradicts the predicted transfer in 3-trials. Although a within-subject comparison of transfer across trial counts (e.g., via repeated-measures ANOVA on JSD values or permutation z-scores) could have been conducted, the absence of significant effects in any condition made such contrasts uninformative. Future work with greater sensitivity or different transfer tasks may help test this hypothesis more conclusively.

This absence of transfer may be explained by the timing and nature of synergy consolidation. Study 1 showed that motor synergy distributions stabilised by Trial 12, suggesting a transition from flexible to more rigid motor control patterns even within short learning windows. Prior research supports this rapid transition: Latash (2010) outlines early-stage motor learning as a period of synergy formation followed by rapid consolidation; Komar et al. (2023) report that solution spaces narrow during early practice, limiting behavioural variability; and Cheung et al. (2009) and Kaufmann et al. (2024) show that synergies stabilise with increased proficiency, reducing trial-to-trial variability and overlap. These findings help explain why longer practice durations might hinder transfer due to increased task-specificity and reduced flexibility. Conversely, the null result in the 3-trial condition suggests that even brief exposure may not provide sufficient consolidation to form transferable synergies, or that task complexity interrupts their application.

Turning our attention to population metrics like velocity, acceleration, and

overlap provided further insights into motor synergy differences between 3-element and 5-element sequences. Velocity consistently showed higher values in the last trial of 3-element sequences compared to the first trial of 5-element sequences, indicating that participants moved faster in simpler tasks. Acceleration differences, though significant, were smaller in effect size, suggesting modest adjustments in movement initiation. Overlap data revealed the most substantial difference, with higher values in 3-element sequences, reflecting greater temporal overlap in participants' actions. These population-level results highlight a shift in motor control strategies when participants transitioned from simpler to more complex sequences, aligning with the individual findings from the rotations.

The cases of disagreement between transition speeds and pose time series data highlight instances where skill learning transfer - suggested by stable transition speeds - did not align with significant changes in motor synergy use. In these specific trial pairs, although transition speeds between the final trials of 3-element sequences and the initial trials of paired 5-element sequences were not significantly different, the pose time series data and JSD analyses revealed significant differences in motor synergy distributions. For example, in some time series comparisons, distinct shifts in finger movement patterns emerged, illustrating that participants generated new motor patterns for more complex sequences. These cases underscore the value of examining nuanced motor control changes that are not reflected in transition speed alone. Thus, even when transition speeds suggested skill transfer, the shifts in pose data and motor synergy distributions indicate that the synergies formed in simpler sequences did not fully transfer to more complex tasks.

Together, these findings provide compelling evidence that motor synergies developed in one skill do not automatically transfer to subsequent skills, especially when task complexity increases. While transition speed improvements suggest that participants learn to perform sequences more fluently, JSD and kinematic differences demonstrate that underlying motor synergy use shifts as tasks become more complex. Notably, cases of disagreement between transition speeds and JSD values were rare, occurring in only 4 out of the 36 total sequence pairs across all rotations,

underscoring that these were specific, rather than pervasive, instances. Additionally, it is important to note that there was no instance where observed JSD was non-significant while transition speeds were significant, further highlighting the robustness of JSD in detecting subtle but meaningful changes in motor synergies. This divergence emphasizes the limitations of relying solely on timing or speed metrics to assess motor skill transfer and underscores the importance of considering both temporal and kinematic dimensions in understanding motor learning.

5.4 Discussion

Study 3 examined the transfer of motor synergies between simpler and more complex motor sequences, focusing on whether practice duration and task structure influence synergy transfer. Participants practised 3-element sequences followed by paired 5-element sequences that embedded the initial sequence, allowing us to evaluate whether synergies developed during the simpler task transferred to the more complex one. While participants exhibited improvements in correct keypresses per second (ckp/s) and transition speeds across trials, distributions of motor synergies used did not transfer between tasks, as quantified with Jensen-Shannon divergence (JSD). Significant differences in synergy distributions between the final trial of 3-element sequences and the initial trial of 5-element sequences suggested that motor synergy transformation is required when transitioning to more complex tasks. These findings highlight the hierarchical and task-specific nature of motor synergies, where the modular control strategies developed for one task may not generalise effectively to tasks with increased complexity (Latash, 2021; Bizzi and Cheung, 2013). At the population level, metrics such as velocity, acceleration, and overlap revealed additional insights into the differences between 3-element and 5-element sequences. Higher velocities and overlaps in 3-element sequences indicated more efficient and temporally coordinated movements in simpler tasks, whereas reduced values in 5-element sequences reflected the increased cognitive and motor demands of task complexity. Further, UMAP with Hamming distance and recursive HDBSCAN clustering allowed us to identify group-level patterns in motor synergy

development, providing a detailed view of how participants adjusted their motor strategies across sequence lengths in spite of shared sequential components. Despite individual variability, the overall results consistently showed that while simpler motor sequences enabled participants to develop modular control strategies, those strategies used for the 3-element sequences required reconfiguration for successful execution of the same sequence once embedded into the 5-element structure.

5.4.1 In the Context of Generalisation Literature

Study 3 results emphasise the utility of JSD in capturing nuanced changes in motor control that are not reflected in traditional performance metrics like correct keypresses per second (ckp/s) or transition speeds. While these measures indicated improvements in execution, JSD analyses revealed significant shifts in motor synergy patterns, showing a reconfiguration of motor commands in response to increased task complexity. This reconfiguration underscores findings in prior research that motor skill transfer is frequently task-specific and context-dependent (Schmidt, 1975; Krakauer and Shadmehr, 2006). Such specificity highlights the challenges of generalising synergies across tasks that differ in structure or complexity.

The lack of synergy transfer observed in this study agrees with adaptation studies that suggest transfer is most effective when tasks share spatial, temporal, and biomechanical characteristics (Wang and Sainburg, 2007). For example, visuomotor adaptation research has consistently demonstrated that generalisation is strongest when trained and transfer tasks are highly similar in workspace, movement dynamics, or sensory feedback (Seidler et al., 2010). In contrast, tasks with increased complexity or altered demands often necessitate significant reorganisation of motor synergies, as observed in Study 3. The hierarchical structure of the paired 3- and 5-element sequences likely constrained the direct transfer of synergies, requiring participants to adjust their motor control strategies for the more complex task. This may have been due to the fact that the embedded 3-element sequences went unrecognised, as the cognitive or hierarchical cue for their execution was misaligned within the 5-element task. Focused on rapid adaptation to meet the speed and accuracy demands of the task, participants may have transformed their syn-

ergies to accommodate immediate needs rather than relying on previously formed patterns. This finding aligns with theories that synergies operate as modular units optimised for specific task demands, where their flexibility diminishes as complexity increases (Tresch and Jarc, 2009; Santello et al., 2016).

At the same time, these results challenge the broader notion of synergy flexibility as described in general motor control theories. While synergies are often characterised as adaptable and capable of supporting novel task demands (Latash, 2021; Bizzi and Cheung, 2013), the observed need for synergy reconfiguration suggests that adaptability may be limited by the degree of overlap between sequential components or task goals. Study 3 findings suggest a more nuanced perspective: although synergies may initially exhibit flexibility, their task-specific refinement during practice may constrain their ability to generalise to tasks that embed those same components in a new structural framework.

The hierarchical structure of motor sequences explored in this study also contrasts with findings from studies on continuous movement tasks. Continuous tasks, such as reaching or walking, often show partial transfer of synergies across different contexts due to the overlapping biomechanical and neural requirements (Krakauer, 2009). However, the motor sequences here likely introduced unique cognitive and motor challenges, particularly as participants transitioned from 3-element sequences to those sequences embedded into the 5-element sequences. This finding extends prior work by illustrating how the nature of real-world motor skills can further constrain generalisation.

In some cases, results observed in Study 3 resemble findings from rehabilitation research, where transfer between trained and untrained tasks is often limited despite functional improvements in related behaviours (Lang et al., 2013; Wolf et al., 2006). Such research highlights that while general motor improvements may be evident, task-specific adaptations frequently require additional practice to consolidate. Study 3 similarly demonstrates that motor synergies do not automatically transfer across tasks, even when they share overlapping components, emphasising the critical role of task structure and complexity in generalisation.

Together, these findings refine existing theories of motor generalisation by showing that synergy adaptability is constrained not only by task demands but also by structural embedding and sequence context. Rather than reusing previously acquired motor modules, participants appeared to reconfigure their synergies in response to the added complexity of the 5-element sequences - even when the underlying components remained the same. These results suggest that generalisation is not simply a matter of shared biomechanical components, but is also shaped by the broader task architecture and cognitive framing. This distinction has methodological implications as well, motivating the group-level clustering and pose-derived metrics used in Study 3 to detect shifts in motor organisation across structurally related but functionally distinct tasks.

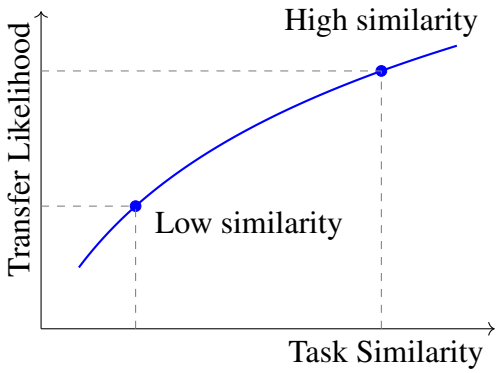

Figure 5.17: Illustration of the hypothesised relationship between task similarity and synergy transfer likelihood. As tasks become more similar in spatial, temporal, and biomechanical features, the probability of reusing pre-formed synergies increases. Transfer is limited when tasks differ structurally or cognitively.

5.4.2 Notes on Group-Level Analyses

Unlike Studies 1 and 2, which focused on individual-level analyses, Study 3 employed group-level methods to accommodate the larger dataset and diverse trial structures. This shift required methodological adaptations, such as the use of UMAP with Hamming distance and recursive HDBSCAN clustering, to identify shared patterns across participants. While these techniques effectively captured group-level trends, they also introduced challenges in interpreting individual variability. For instance, differences in learning rates and motor control strategies

among participants may have contributed to the observed variability in JSD values.

Group-level analyses also highlight the benefits of using metrics like velocity, acceleration, and overlap, to contextualise and explain cluster-based findings in interpretable terms. These measures provided a comprehensive view of motor performance, revealing consistent trends across participants despite individual differences. However, future studies should aim to integrate group- and individual-level analyses to pursue a nuanced understanding of motor skill/synergy transfer.

5.4.3 Study Limitations

In parallel with these modifications for group-level analyses, I must also acknowledge the limitations of study design. The first limitation concerns sequence overlap and repetition. Due to the counterbalanced structure of the study, in which all participants practised six different initial–paired sequence pairs in rotation, each 3-element sequence not only appeared in its paired 5-element sequence but also occasionally reappeared in later 5-element sequences. This partial redundancy may have introduced familiarity effects unrelated to the intended transfer. For instance, sub-sequences such as 241 or 234 were embedded in multiple different 5-element sequences, raising the possibility that participants benefited from general exposure rather than true transfer of a previously formed synergy structure.

Another important limitation is the absence of a no-practice control condition. Without a control group that performs the 5-element sequence without prior exposure to the embedded 3-element sequence, it is difficult to definitively attribute any observed differences in performance or synergy structure to transfer, rather than to general task familiarity or baseline variability. Although my within-subject design permitted testing multiple transfer scenarios under different levels of prior practice, it did not allow us to quantify the added benefit (or lack thereof) of prior training on the embedded 3-element subsequence.

This trade-off between ecological realism and experimental control is a common challenge in motor learning research. The rotation-based design reduced ordering effects and increased sequence diversity, but limited the interpretability of any one trial pair. Future studies could adopt a hybrid approach - retaining the

sequence-pair structure while adding a between-subjects control group - to better isolate the causal contribution of prior practice to motor synergy transfer.

5.4.4 GoPro Video Challenges

A key challenge in Study 3 was managing video data recorded using GoPro cameras. Due to file size limitations, GoPro cameras automatically split recordings into multiple files. For each participant session, this resulted in video files with identical prefixes but unique suffixes (e.g., `GX010001.MP4` and `GX010002.MP4`). To address this, custom Python scripts were developed to programmatically resize and stitch videos back together.

The solution consisted of two main steps: (1) resizing videos to a consistent resolution (1024x576) to reduce file size and processing overhead, and (2) stitching files with matching prefixes and sequential suffixes into single, continuous videos. Below is the implementation:

```
import os
import re
import numpy as np
from moviepy.editor import VideoFileClip,
    concatenate_videoclips
from send2trash import send2trash
from PIL import Image

# Step 1: Resize videos to reduce resolution
def custom_resize(clip, newsize):
    def resize_image(image):
        pil_image = Image.fromarray(image)
        resized_pil_image = pil_image.resize(newsize, Image.
            LANCZOS)
        return np.array(resized_pil_image)
    return clip.fl_image(resize_image)

def resize_videos(videos, output_directory, resolution):
```

```
    for video_path in videos:
        clip = VideoFileClip(video_path)
        clip_resized = custom_resize(clip, resolution)
        base_filename = os.path.basename(video_path)
        new_filename = f"{os.path.splitext(base_filename)[0]}
            _resized.mp4"
        output_path = os.path.join(output_directory,
            new_filename)
        clip_resized.write_videofile(output_path, codec='
            libx264', audio_codec='aac')
        clip.close()
        clip_resized.close()
        send2trash(video_path)
        print(f"Original file {video_path} moved to trash.
            Resized file saved as {output_path}")

# Step 2: Stitch videos with matching prefixes
def find_videos_with_prefix(directory, prefix):
    videos = []
    pattern = re.compile(rf"^{prefix}.*\.MP4$")
    for filename in os.listdir(directory):
        if pattern.match(filename):
            videos.append(os.path.join(directory, filename))
    return videos

def stitch_videos(video1_path, video2_path, output_path,
    resolution):
    clip1 = VideoFileClip(video1_path)
    clip2 = VideoFileClip(video2_path)
    fps = clip1.fps
    clip1 = custom_resize(clip1, resolution)
    clip2 = custom_resize(clip2, resolution)
```

```
    final_clip = concatenate_videoclips([clip1, clip2])
    final_clip.write_videofile(output_path, fps=fps, codec='
        libx264', audio_codec='aac')
    send2trash(video1_path)
    send2trash(video2_path)
    print(f"Stitched video saved as {output_path}")

# Example usage
directory = "/path/to/videos"
output_directory = "/path/to/output"
prefix = "GX01"
resolution = (1024, 576) # Desired resolution

# Resize and stitch videos
videos = find_videos_with_prefix(directory, prefix)
if len(videos) > 1:
    for i in range(0, len(videos) - 1, 2):
        video1, video2 = videos[i], videos[i + 1]
        output_path = os.path.join(output_directory, f"
            stitched_{i // 2}.mp4")
        stitch_videos(video1, video2, output_path, resolution)
```

Critical Elements of the Code:

- **Resizing Videos:** The `custom_resize` function ensures all videos are resized to a consistent resolution (1024x576) using Lanczos filtering for high-quality downscaling. This reduces file size and processing overhead while maintaining sufficient visual quality for kinematic analysis.
- **Finding Video Files:** The `find_videos_with_prefix` function uses regular expressions to identify all video files with a specified prefix in a directory. This ensures only relevant files are processed.
- **Stitching Videos:** The `stitch_videos` function concatenates two video

clips, resizing them to the specified resolution and ensuring consistent frame rates. The original files are moved to the trash after successful stitching to avoid duplication.

- **Automation:** The script is designed to handle multiple video pairs automatically, iterating through files with matching prefixes and stitching them into a single output file.

Intention: I include this solution in the discussion only to provide a practical resource for researchers facing similar challenges with GoPro video management. By combining automated resizing and stitching, the script streamlines data processing and integration of video data into analysis workflows.

Chapter 6

GENERAL DISCUSSION

"I may not have gone where I intended to go, but I think I have ended up where I needed to be."

— Douglas Adams

6.1 Thesis Overview

This thesis examined the early stages of motor sequence learning, focusing on the emergence, timing, and development of motor synergies, as well as their transfer between motor sequence skills. Across three experimental studies, I investigated how synergies are initially formed, how they develop during early-stage practice, and whether they transfer between tasks of differing complexity. While the findings provide insights into these processes, they also reveal the limits of transfer, suggesting that synergy reconfiguration is often necessary for new task demands. Below, I summarise the key contributions of each chapter, emphasising the broader implications for motor learning research.

6.1.1 Chapter 2: Methods

Chapter 2 detailed the methods developed for Studies 1 and 2, addressing the unique challenges posed by remote data collection during the COVID-19 pandemic. These methods enabled rigorous behavioural quantification and unsupervised classifica-

tion of motor behaviours using tools such as DeepLabCut for pose estimation and HUB-DT for behavioural annotation. I highlighted the integration of UMAP and HDBSCAN for dimensionality reduction and clustering, enabling the identification of motor synergies at the trial level.

6.1.2 Chapter 3: Early Motor Synergy Development

Study 1 investigated the formation and stabilisation of motor synergies during early-stage learning. By analysing synergy distributions across trials, I found that rapid gains in synergy organisation occurred within the first 12 trials, followed by stabilisation as participants approached a performance plateau. These findings highlighted the timeline of synergy development, suggesting that early practice is critical for establishing modular control strategies that underpin skilled performance.

6.1.3 Chapter 4: The Role of Rest in Synergy Development

Study 2 explored the role of rest in motor learning, revealing that offline periods between practice sessions facilitated rapid motor synergy reorganisation. Results demonstrated that rest enhanced skill consolidation, with synergy transformations contributing significantly to performance gains. This study provided evidence for the motor system's capacity to alter control strategies over periods of rest rather than active practice.

6.1.4 Chapter 5: Synergy Transfer Across Tasks

Study 3 examined whether motor synergies developed during simpler tasks could transfer to more complex ones. Results showed limited direct transfer, as participants required reconfiguration of synergies when transitioning to tasks with increased complexity. The findings underscored the hierarchical and task-specific nature of motor synergies, where previously learned strategies must often be adapted to accommodate new task demands. However, the sensitivity of synergy transfer to training duration and cue placement warrants further investigation, particularly concerning cognitive and contextual factors that may influence transfer efficacy.

6.1.5 Integration with Broader Literature

The findings of this thesis contribute to a growing body of research on motor sequence learning by providing a detailed examination of the early stages of motor synergy formation, development, and transfer. Prior studies have established that motor synergies serve as modular units of control, enabling the central nervous system to simplify the coordination of complex movements (Bizzi and Cheung, 2013; Latash, 2021). However, much of this work has focused on well-learned behaviours or adaptation processes, with limited emphasis on how synergies emerge and stabilise during the early stages of learning a continuous motor sequence skill. This thesis addresses this gap by highlighting the timing and rapid organisation of motor synergies during initial practice, as seen in Study 1, and the critical role of rest in facilitating their transformation and consolidation, as demonstrated in Study 2.

The negligible transfer of synergies in Study 3 aligns with prior research on the context-specific nature of motor learning (Schmidt, 1975; Seidler et al., 2010). Previous studies on visuomotor adaptation have shown that transfer is most effective when tasks share spatial and biomechanical similarities, underscoring the challenges of applying learned motor strategies to tasks with altered cues or demands (Wang and Sainburg, 2007). The hierarchical structure and embedded components of the sequences expand on this work, revealing that even minor cognitive or structural modifications can disrupt synergy transfer. This insight complements the growing recognition of task-specific learning mechanisms, where synergies reflect an optimisation tailored to the immediate demands of the practised task.

Moreover, use of unsupervised clustering and dimensionality reduction techniques, such as HDBSCAN and UMAP, extends existing approaches to behavioural analysis by offering data-driven, non-linear methodologies for quantifying synergy formation. These methods align with recent trends in motor neuroscience, where advanced computational tools are increasingly employed to uncover latent patterns in high-dimensional movement data (Hausmann et al., 2021; Wagner et al., 2020). The use of JSD to capture changes in synergy distributions further shows how metrics can deepen our appreciation for motor control, bringing together traditional

kinematic analyses and modern data science approaches.

In synthesising these findings with existing literature, this thesis contributes to an understanding of motor sequence learning. While supporting the modular and hierarchical framework of motor synergies, it also raises questions about their flexibility and the extent to which early-stage synergies can be repurposed for tasks of greater complexity. These contributions provide an interesting waypoint for future investigations into the interplay between task structure, cognitive demands, and motor control strategies in both typical and impaired populations.

6.1.6 Practical Implications and Future Directions

The findings of this thesis have possible practical implications for understanding and enhancing motor sequence learning, particularly in educational, clinical, and rehabilitative contexts. The identification of distinct stages in motor synergy development, as seen in Studies 1 and 2, underscores the importance of tailoring training protocols to the specific phases of skill acquisition. Early-stage learning, characterised by rapid synergy formation, may benefit from intensive, focused practice sessions that prioritise consistency and accuracy. Conversely, the role of rest in facilitating skill consolidation, as demonstrated in Study 2, highlights the need to balance practice with sufficient offline periods to optimise learning outcomes.

From a clinical perspective, Study 3 results raise important considerations for rehabilitation. The need for synergy reconfiguration when transitioning to more complex tasks suggests that interventions should carefully sequence tasks to build upon previously learned skills. Incorporating hierarchical task structures, where simpler components are embedded saliently within progressively complex activities, may help patients incorporate motor strategies more effectively. This approach could be beneficial for individuals recovering from neurological injuries, such as stroke, where supplementing existing clinical synergies with biomechanically healthy, or typical,motor synergies is a primary goal.

The application of computational tools, such as UMAP and HDBSCAN, further extends the scope of motor learning research by enabling detailed, unsupervised analysis of high-dimensional behavioural data. These methods are a framework for

automated systems capable of monitoring skill acquisition and offering feedback to learners and clinicians in real time. Such technologies could be integrated into wearable devices or mobile applications, making them accessible to a wide range of users. Moreover, the use of metrics like Jensen-Shannon divergence to quantify changes in motor synergies presents an opportunity to refine assessments of motor function, particularly in clinical populations where traditional measures may be insufficient to capture subtle behavioural adaptations.

Looking forward, several questions remain unanswered and warrant further exploration. For example, how do cognitive and contextual factors, such as cue placement and task familiarity, influence the transfer of motor synergies? Additionally, future work could test my findings across diverse populations, including individuals with motor impairments, to understand a broader applicability.

Finally, this thesis supports the need for interdisciplinary collaboration in motor learning research. Bridging neuroscience, data science, and clinical practice can accelerate the development of innovative tools and interventions that enhance our understanding of motor control while addressing practical challenges in skill acquisition and rehabilitation. By integrating empirical findings with cutting-edge technologies, future research has the potential to transform both theoretical perspectives and real-world applications in motor learning.

6.2 Considerations for the Collection of Remote Data

In Chapter 2 of this thesis, I provided a comprehensive guide to the methods employed in Studies 1 and 2, both of which utilised remote data collection technologies and procedures. As noted, my remote methods were designed to meet the challenges of studying motor skill learning behaviour whilst stay-at-home orders were in effect in the United States during the COVID-19 pandemic. Indeed, the initial aim of this doctoral research was to work directly with stroke patients to develop and validate an in-home clinical evaluation system, using video to track specific upper limb movements and analyse data to quantify recovery. However, the COVID-19 pandemic and resulting lockdowns imposed significant restrictions on

in-person research activities, limiting my ability to engage with this patient population as planned. Despite these challenges, working with the participants we were able to include provided valuable insights into the practicalities of implementing a process for remote data collection for patients in the home setting. These experiences highlighted key considerations in both hardware and software domains, informing strategies to scale the project effectively while ensuring accessibility.

Hardware Considerations

Implementing an at-home clinical evaluation system for stroke patients at scale necessitates careful planning of hardware requirements to ensure effective data collection and usability across diverse populations.

Device Selection and Variability

The first hardware to consider is a camera capable of capturing video suitable for pose estimation. To collect data from a large number of patients in their homes, the system must utilise devices that are widely accessible and affordable across socioeconomic statuses. One practical solution is leveraging the ubiquity of smartphones or tablets, which often come equipped with advanced camera technology. These devices could capture video data without imposing additional costs on patients. Developing a user-friendly application compatible with a range of devices and operating systems could facilitate adoption. However, this approach introduces variability in hardware capabilities, such as differences in camera resolution, frame rates, processing power, and sensor quality, which can affect data consistency and quality. Standardising data collection protocols is essential to mitigate these issues. This could involve setting minimum hardware requirements for device compatibility or implementing software-based adjustments to account for hardware differences.

Ease of Use and Accessibility

Ensuring ease of use for patients who may have limited technical proficiency or physical impairments is important. The hardware setup should require minimal configuration and intuitive interfaces with clear visual cues. Voice command or simplified touch support may also enhance usability for a broader range of patients.

Connectivity and Data Transfer

Connectivity is another significant consideration. Patients may have limited access to high-speed internet, which poses challenges for transmitting video data. Implementing data compression techniques, such as efficient video codecs or transmitting only essential data (e.g., extracted pose coordinates instead of raw video), can reduce bandwidth requirements. Alternatively, allowing for offline data collection with periodic uploads when connectivity is available can accommodate patients with intermittent internet access.

Environmental Factors and Setup Guidance

Environmental factors within patients' homes must also be considered. Variations in lighting, background, camera positioning, and available space can all impact pose estimation or patient performance. Algorithms must therefore be designed to tolerate these challenges, possibly through data augmentation or preprocessing techniques. Providing patients with guidance on optimal setup is also important. This could include recommendations on camera placement, suggestions for suitable locations within the home (e.g., areas with adequate lighting), and instructions to minimise occlusions (e.g., avoid loose clothing that obscures body movements). Including real-time feedback mechanisms within the application can help patients adjust their setup if the system detects suboptimal conditions.

Software Considerations

Scaling relevant project software involves developing systems capable of processing and analysing large volumes of data efficiently and accurately.

Scalability and Cloud Infrastructure

As the number of patients increases, the software must handle extensive datasets without compromising performance or responsiveness. Implementing scalable algorithms and leveraging distributed computing resources are key strategies to achieve this goal. Any adapted self-supervised or unsupervised framework for time series data will require significant computational power for training and inference. Utilising cloud computing platforms can provide the necessary resources, offering

scalability through on-demand provisioning of computational and storage capacities. Cloud services, such as Amazon Web Services (AWS), Google Cloud Platform (GCP), or Microsoft Azure, offer tools for distributed computing, allowing the system to parallelise processing tasks across multiple servers. This reduces processing times and enables the handling of large datasets efficiently.

Automated Pipelines and Data Flow

Automated data processing pipelines are essential for handling continuous data streams from patients. These pipelines should encompass data ingestion, preprocessing, feature extraction, model training, and real-time inference. Automation minimises the need for manual intervention, reducing the potential for errors and improving efficiency. Implementing technologies like Apache Kafka or AWS Kinesis can facilitate real-time data streaming and processing.

Storage Solutions and Data Management

Developing scalable databases and storage solutions capable of handling time series data effectively is another important consideration. Employing databases optimised for high write and read throughput, such as time series databases (e.g., InfluxDB, TimescaleDB) or NoSQL databases (e.g., MongoDB, Cassandra), can improve data retrieval and storage efficiency. Data indexing, sharding, and partitioning strategies can enhance performance for large datasets by distributing the data across multiple nodes and facilitating parallel access.

Model Maintenance and Monitoring

The software must also include mechanisms for model updates and maintenance. As more data becomes available, models may need retraining to improve accuracy and generalisability. Implementing continuous integration and continuous deployment (CI/CD) pipelines allows for seamless updates to the system without disrupting service. Tools like Jenkins, GitLab CI/CD, or AWS CodePipeline can automate the build, test, and deployment processes. Monitoring model performance in production is essential to detect model drift or degradation in accuracy. Implementing monitoring tools and alerting systems can help maintain model reliability. Tech-

niques like A/B testing or canary deployments can be used to evaluate new models before full-scale deployment.

Scaling an at-home clinical evaluation system involves leveraging available hardware, ensuring the system is accessible, and reliably collecting diverse data from patients. On the software side, employing scalable architectures and robust data processing pipelines enables efficient handling of datasets and timely insights. Addressing these considerations are a start for future developments in remote patient monitoring and personalised rehabilitation.

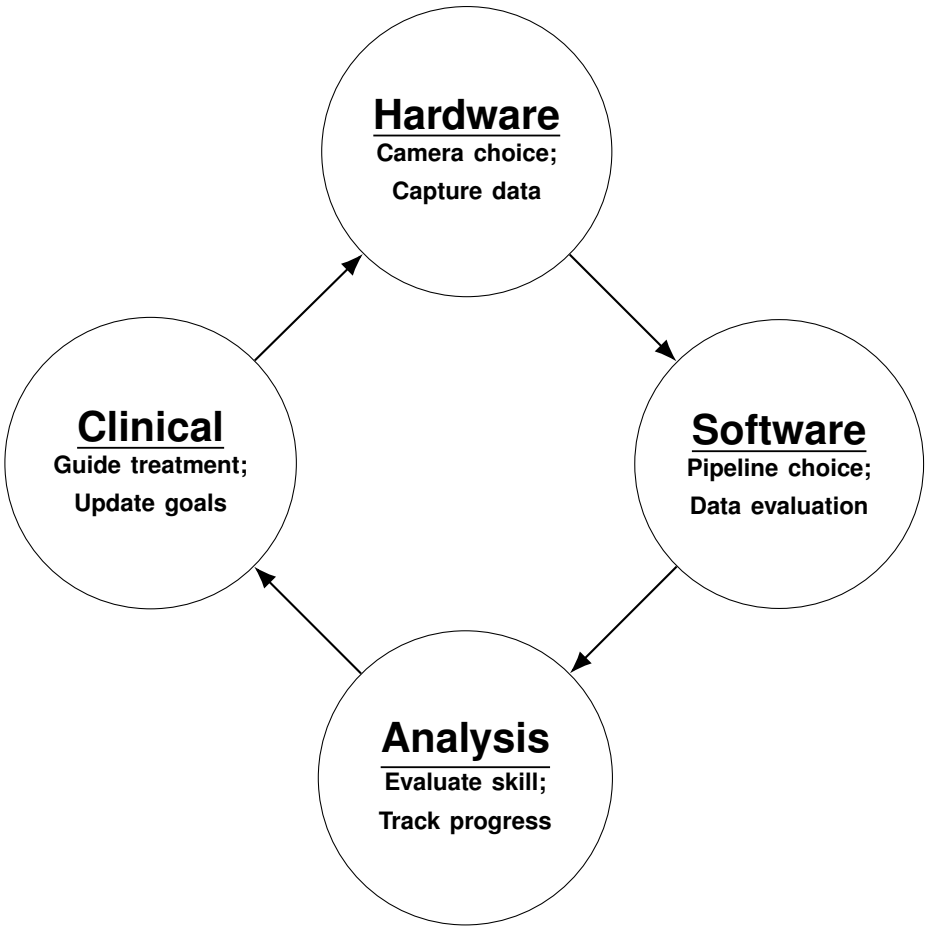

Figure 6.1: Circular pipeline for remote data collection and evaluation. Hardware (e.g., cameras) enables data capture, which is processed by Software pipelines, analysed to track skill learning and motor recovery, and fed back into Clinical decision-making loops.

6.3 Synergy and Skill Requirement

In this section, I consider how the requirements of a task drive the emergence and configuration of motor synergies. Findings across Studies 1–3 suggest that the synergies observed were not solely the result of biomechanical or neural constraints, but also significantly shaped by task demands - specifically, the instruction to perform each sequence as quickly and accurately as possible. This constraint consistently directed participants toward movement strategies optimised for speed-accuracy trade-offs, rather than theoretical efficiency or exploratory movement diversity.

This interpretation aligns with prior work showing that synergies are flexible, task-dependent constructs. Latash (2021) describes synergies as modular systems capable of adaptation even under tight constraints. Similarly, Bizzi and Cheung (2013) demonstrate that synergy composition reflects the demands of the environment and goal. The speed–accuracy requirement in my studies likely constrained the solution space, encouraging the consolidation of highly repeatable movements.

Study 3 offers a useful lens through which to examine this constraint-driven optimisation. As participants transitioned from 3-element to 5-element sequences, I observed that synergies often reorganised entirely - despite the 3-element sequence being embedded within the longer, complex task. This reconfiguration may have reflected the need to maintain speed and accuracy over a more complex structure, highlighting how increased task demands reshape motor control.

To illustrate how alternative constraints might yield different outcomes, consider a hypothetical version of the same task in which participants are instructed to prioritise variability or novelty of movement, rather than speed. Under such instructions, one might expect greater within-participant variability and slower convergence toward stable synergies, or even the emergence of different exploratory strategies altogether. Alternatively, if participants were explicitly instructed to maintain uniform timing between keypresses or emphasise movement smoothness, the resulting synergy patterns could reflect different coordination principles, such as tempo regulation or biomechanical efficiency rather than pure speed.

This possibility raises a broader question about the boundaries of synergy adaptability: how plastic are these control strategies when the task goal shifts? Al-

though this thesis focused on speed-constrained learning, future studies could investigate the reconfiguration of motor synergies under alternate performance goals. Such work may help distinguish between synergies that are robust across contexts and those that are tightly coupled to a particular task structure.

Finally, this discussion complements findings in motor adaptation, where task goals have long been recognised as a key determinant of learning trajectories (Seidler et al., 2010; Wang and Sainburg, 2007). Across domains, task specificity appears to drive not only the acquisition of skill but also the architecture of the motor solutions used to support it. The findings in this thesis contribute to that body of work, offering new insights into the relationship between motor learning, task constraints, and synergy flexibility.

6.4 Synergy: An Ill-Defined Problem

Despite the widespread use of the term *synergy* in motor control research, its definition and application remain highly variable and context-dependent. In this thesis, I examined motor synergies through a specific lens: as clusters of motor behaviours identified using unsupervised computational techniques, such as UMAP and HDBSCAN. While this approach yielded valuable insights into the timing, emergence, and development of synergies, it also raised questions about what *synergy* truly signifies in the broader context of motor learning.

Latash (2021) aptly summarises this challenge, describing synergy as a concept with "many solutions." Synergies may be defined and calculated based on muscle activation patterns, joint kinematics, or even neural activity, depending on the experimental framework. This diversity of definitions is both a strength and a limitation for the field. On one hand, it reflects the multifaceted nature of motor control; on the other, it complicates cross-study comparisons and theoretical integration.

The issue is further exacerbated by the validation loops inherent in synergy research. Researchers often define synergies based on specific tasks and methods, then validate their findings using metrics optimised for those same tasks. While this approach ensures internal consistency, it risks creating isolated pockets of knowl-

edge rather than a unified understanding of motor control or motor synergy. For instance, studies focusing on muscle synergies often highlight their modularity as evidence of neural optimisation (Bizzi and Cheung, 2013), while kinematic analyses may emphasise biomechanical efficiency (Tresch and Jarc, 2009).

This lack of consensus prompts a critical question: What are researchers truly pursuing when they study synergy? Is it a purely descriptive construct, or does it represent an underlying principle of motor organisation? Addressing these questions requires interdisciplinary collaboration and methodological standardisation, drawing from neuroscience, biomechanics, and computational modelling.

Future research must also consider the contextual dependency of synergy formation. For example, the synergies observed in this thesis were heavily influenced by the requirement to perform tasks rapidly and accurately. In other contexts, such as rehabilitation or creative motor tasks, entirely different synergy configurations may emerge. Exploring these variations could help clarify whether synergies represent universal principles of motor control or task-specific adaptations.

Domain	**Definition of Synergy**	**Measurement Modality**
Muscle Activation	Synergies as coordinated patterns of muscle co-activation that reduce the dimensionality of control	Electromyography (EMG)
Joint Kinematics	Synergies as covarying joint angles or limb segment trajectories during movement	Motion capture; Inertial measurement units (IMUs)
Neural Activity	Synergies as structured patterns of population-level neural firing or cortical organisation	fMRI, EEG, or invasive recording
Behavioural Clustering	Synergies as recurring clusters of behavioural states or movement patterns identified using unsupervised learning	Pose estimation + UMAP/HDBSCAN or other clustering algorithms

Table 6.1: Summary of synergy definitions and corresponding measurement modalities across domains.

Ultimately, synergy remains an ill-defined but indispensable concept in motor

control research. By recognising its limitations and embracing its diversity, the field can move toward a more comprehensive and integrative understanding of how humans and other animals achieve skilled movement.

6.5 Neurology of Rest-Based Calculations

Rest plays a critical role in motor skill learning, acting as a period for the consolidation of both spatial-ordinal information (the *what to do* in a sequence) and sensorimotor strategies (the *how to do it*). This thesis demonstrated how the the motor system leverages brief periods of rest to reorganise and refine motor synergies on a trial-by-trial basis. These findings underscore a capacity to rapidly update cognitive representations of task demands and concurrently optimise motor behaviours.

The central nervous system's ability to execute complex operations within seconds is of note. It involves integrating feedback from prior movements, predicting future task requirements, and restructuring motor synergy modules to accommodate evolving demands. This rapid consolidation highlights the adaptability of motor circuits, which can process and encode a high volume of sensorimotor information even during short pauses in task performance.

While prior studies have emphasised the role of extended rest periods in offline skill consolidation, Study 2 findings support modern evidence that learning can occur within seconds. Interleaved rest intervals facilitate iterative updates, allowing the motor system to reorganise control strategies. This is distinct from online learning, where changes occur during active task execution and are often driven by error correction and feedback-based adaptation. In contrast, rest-based reorganisation appears to reflect a process that coordinates neural resources across networks to solidify and reconfigure synergies without concurrent motor output.

Emerging work by Cohen, Buch, and Bonstrup supports this distinction, showing increased functional connectivity between motor, sensory, and prefrontal networks during rest. These findings indicate a heightened state of coordination and information exchange during brief pauses, underscoring the importance of short, strategically placed rest intervals in skill acquisition (Buch et al., 2021).

Given these findings, future research should delve deeper into the neurological underpinnings of rest-based learning, exploring how motor circuits dynamically balance the consolidation of spatial-ordinal and sensorimotor information. Key questions include whether the benefits of rest extend to more complex motor tasks, how rest duration modulates synergy reorganisation, and how rest-based consolidation differs in populations with varying neurological profiles.

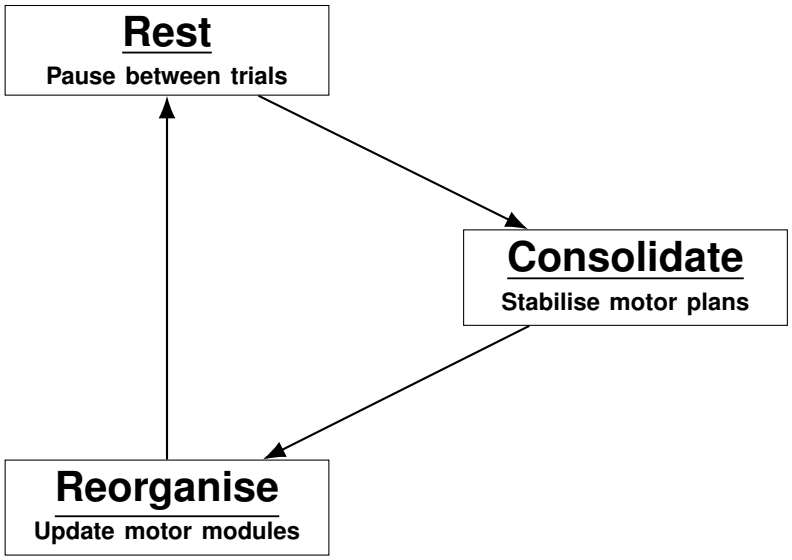

Figure 6.2: Schematic of the rest → consolidation → synergy reorganisation cycle in Study 2. Rest intervals trigger internal processing that consolidates motor experiences, which then facilitates reorganisation of motor synergies in preparation for subsequent trials.

6.6 Rest in Clinical Participants

Assuming a critical stance on the topic of rest and the results of Study 2, I turn briefly to a conversation on the nature of rest in clinical populations and how extant results in healthy participants may not reasonably replicate in those with neurological insult. Resting-state neurological activity in individuals recovering from stroke - particularly in the subacute and chronic phases - differs significantly from that of healthy populations. These differences are often linked to disrupted connectivity in motor and sensory networks, which play central roles in skill consolidation. Studies have shown that stroke patients exhibit reduced functional connectivity between primary motor areas and distributed networks important for learning, potentially limiting their capacity to benefit from rest-dependent reorganisation processes (Grefkes

and Fink, 2011; Stagg and Johansen-Berg, 2013).

The findings presented in this thesis raise important questions about the generalisability of rest-based learning paradigms to such populations. While healthy participants demonstrated rapid synergy reorganisation following brief rest intervals, it remains unclear whether individuals with neurological injuries - especially those in the chronic stage - experience similar benefits. Subacute stroke patients may retain a greater degree of neuroplastic potential, but they too show signs of altered resting-state dynamics. For instance, hyperexcitability in motor circuits, frequently observed post-stroke, often results from imbalances in interhemispheric inhibition and reduced cortical GABAergic activity (Stinear et al., 2015). This pathological state may interfere with the consolidation processes observed in healthy adults.

A critical future direction involves piloting the same sequence learning paradigm used in this thesis - particularly Study 2 - in clinical populations using the existing keyboard interface. This low-cost, at-home compatible device could be used to assess whether rest-related synergy reorganisation is observable in post-stroke individuals. By targeting both subacute and chronic stroke patients, future work can better characterise how the time since injury influences rest-state consolidation. Rest intervals could be manipulated to examine dose-responsiveness, while neuroimaging or portable EEG could provide insight into the neural substrates of rest-related changes. Such studies would help determine whether rest retains the same functional role in clinical contexts or whether additional interventions are needed to unlock its rehabilitative potential. Understanding these differences is vital for translating synergy-based frameworks to rehabilitation, particularly for developing scalable, rest-enhanced training protocols suited to home-based recovery.

6.7 Defining Average Movement

Remaining on the topic of motor skill learning research as a possible benefit to clinical populations, I am reminded of another critical idea in the field that is difficult to approach, i.e. agree on terms for, and even more difficult to solve: defining “average” movement. The concept of average movement is a central challenge for motor

impairment, neurorehabilitation and the computational evaluation of movement and recovery following stroke. While the methods developed in this thesis offer tools for analysing motor development using unsupervised machine learning techniques, the question of what constitutes "average movement" remains inherently complex. The complexity is rooted in inter-individual variability in motor strategies, biomechanics, and demographics. How we solve for that variability to capture a true and useful metric for average movement remains elusive.

Work on motor control and rehabilitation highlights the variability of motor strategies and behaviour within tasks. For example, in an exploration of compensatory strategies in individuals with stroke, Cirstea and Levin (2000) and Levin et al. (2002) observed unique adjustments, such as increased trunk involvement during reaching tasks, that enabled functional goal achievement despite motor impairments. These findings emphasise that "typical" movement is not a fixed standard but rather a dynamic construct that varies across individuals and contexts.

In clinical practice, frameworks such as the Fugl-Meyer Assessment of Motor Recovery (FMA) and the NIH Stroke Scale (NIHSS) are used to approximate functional impairment and recovery benchmarks. The FMA, for instance, provides a quantitative index of motor function for the upper extremity, reflecting reflexes, coordination, and movement quality. The NIHSS includes upper limb motor components within a broader neurological examination. While both tools offer clinically actionable thresholds and serve as proxies for "average" impairment levels, they do not define biomechanical averages or movement profiles. Instead, they reduce complex motor behaviour to ordinal scores, which are useful for stratifying patients but insufficient for modelling motor strategy variation at the kinematic level.

From a clinical perspective, the lack of pre-stroke movement data for most patients further complicates the establishment of a baseline for recovery. As clinicians often rely on relative improvements within each patient rather than comparisons to a universal standard, the absence of a universally accepted definition of "average" movement raises questions about the validity and reliability of cross-patient assessments. Levin and Demers (2021) have argued for personalised rehabilita-

tion approaches that consider each patient's unique motor patterns, a perspective that aligns with the notion that motor recovery should account for biomechanical idiosyncrasies rather than attempting to standardise movement across populations.

Research conducted at National Hospital for Neurology and Neurosurgery (NHNN) also highlights the individualised nature of motor recovery, emphasising that rehabilitation must account for the unique neural and functional profiles of each patient (Ward et al., 2019). Work on intensive upper limb neurorehabilitation demonstrates that tailored, context-specific programmes can lead to significant functional improvements even years post-stroke, challenging the notion of a uniform "typical" movement pattern. This perspective is crucial when applying clustering techniques like UMAP and HDBSCAN to create population-level maps of motor behaviours, as these tools, while powerful, risk oversimplifying the diversity of human movement. By integrating the insights of clinical research, future applications of modern computational methods may better supplement rehabilitation practises and capture the nuanced, individualised nature of motor strategies.

Future research could aim to address these challenges by expanding the scope of movement datasets to include more diverse populations, tasks, and contexts. Longitudinal studies that track motor behaviours over time could help identify common trajectories of recovery while accounting for individual variability. Additionally, incorporating insights from work on compensatory strategies and personalised rehabilitation could refine clustering methods to better capture the dynamic and context-specific nature of motor control. In short, defining "average" movement remains a complex and critical question that demands the integration of individual variability, contextual factors, and broader theoretical frameworks.

6.8 Concluding Remarks

This thesis furthers our understanding of motor sequence learning by investigating the emergence, timing, and development of motor synergies, as well as their transfer across tasks of differing complexity. Through three studies, I demonstrated the central role of task structure, rest, and practice duration in shaping motor synergies

during early-stage learning. Additionally, the methods developed for behavioural quantification and unsupervised classification offer scalable tools for analysing motor behaviours in both research and clinical settings.

These findings have implications for motor learning and rehabilitation, particularly in understanding how synergies evolve to meet the demands of specific tasks. Results emphasise the mutable nature of the motor system, which not only consolidates spatial-ordinal and sensorimotor information with remarkable efficiency but also reorganises motor strategies for each challenge. By addressing key questions about the transferability and specificity of motor synergies, this work contributes to ongoing discussions about the modular and hierarchical nature of motor control.

In closing, this thesis represents a contribution to the study of motor sequence learning and motor synergy, highlighting both the capabilities and complexities of the human motor system. I explored the interplay between motor skill sequence learning, practice, rest, and synergy development, and any insights gained here reflect the collective and collaborative efforts of a scientific community dedicated to understanding human learning and its applications by way of neuroscience, data science, and clinical expertise. This research would not have been possible without each of these fields and I believe any successes are owed to working together.

* * *

Anyway. We've since moved on from the Paris Olympics and are currently immersed in my daughter's ballet class at the Royal Academy of Dance. I watch from the house seats - she moves across the floor with an intensity and joy that feels larger than her small frame. Each plié and pirouette is a puzzle. How does she coordinate her movements so effortlessly, finding balance and fluidity in something so unfamiliar? She's only just learned these steps, yet her body seems to know what to do. It is fascinating, and the scientist in me itches to understand the mechanics that make it all possible. But then I hear her laugh when she misses a turn, resetting with a determination that's all her own. And in that moment, it doesn't matter if I know. Instead, I just sit, watch, and wonder - how does she do that?

Appendix A

SELF-SUPERVISED LEARNING FOR CLINICAL PHENOTYPING OF POSE DATA

A.1 Rationale

Stroke remains one of the leading causes of death and long-term disability worldwide, with an estimated 13.7 million new cases occurring each year globally (Feigin et al., 2021). Survivors often endure a spectrum of motor, sensory, and cognitive impairments, the severity and specific manifestations of which vary significantly between individuals. The complexity of motor impairments post-stroke poses significant challenges; even subtle differences in movement patterns can profoundly affect a patient's ability to perform daily activities and respond to rehabilitation. Traditional clinical assessments may not fully capture these nuanced differences, leading to generalized treatment plans that may not address individual needs effectively. Therefore, machine learning-based phenotyping of post-stroke symptoms may offer a more nuanced understanding of the specific patterns and complexities of motor impairment that clinicians can use to tailor rehabilitation strategies.

A.2 The Need for Democratizing At-Home Clinical Evaluations

Socioeconomic barriers can significantly impede access to consistent, high-quality rehabilitation for stroke survivors. Patients from lower socioeconomic backgrounds often face challenges in accessing care due to limited resources, geographic isolation, or the inability to attend frequent in-clinic evaluations. At-home clinical evaluation systems offer a promising solution to these barriers by ensuring equitable access to care. Moreover, home-based rehabilitation allows for the monitoring of patients in a more naturalistic environment, providing ecologically valid assessments of motor function. The integration of advanced technological solutions, such as computer vision and machine learning, into at-home evaluations can further democratize healthcare by making high-quality clinical tools available at lower costs and with greater frequency. These technological advancements facilitate remote assessments and personalized rehabilitation programs, augmenting clinical ability to monitor the recovery process, adjust treatment plans accordingly, and aim to improve patient outcomes regardless of socioeconomic status.

A.3 Technological Advancements in At-Home Clinical Evaluations

Technological advancements, particularly in the fields of computer vision and machine learning, have opened new possibilities for at-home healthcare evaluations. In recent years, computer vision models have been successfully applied to a variety of healthcare challenges, including gait analysis, motion tracking, and gesture recognition. Specifically, the use of pose estimation algorithms, such as OpenPose (Cao et al., 2021) and MediaPipe (Lugaresi et al., 2019), has demonstrated the potential for precise and non-intrusive tracking of human movement, which is especially pertinent in the assessment of motor function in stroke patients.

Time series analysis plays a crucial role in understanding the dynamics of human motion over time. Advancements in machine learning, including unsupervised

and self-supervised learning techniques, have enabled more refined analyses of sequential patient data without the need for extensive labelled datasets. Models such as SimCLR, originally developed for visual data, can be adapted for time series data to capture temporal patterns in patient movements. By incorporating time series-specific augmentations and architectures suited for sequential data, these models make more capable the analysis of motion patterns over time.

These technological advancements are particularly valuable for large-scale data collection in home settings, where it is often impractical to manually label every instance of patient movement. When combined, computer vision and machine learning offer powerful tools for continuous and detailed tracking of post-stroke motor recovery, providing clinicians with deeper insights into patient recovery.

A.4 Addressing the Problem of Refined Phenotyping with Machine Learning

The application of machine learning algorithms presents a novel approach to refining the phenotyping of post-stroke motor impairments. By adapting learning models like SimCLR to handle time series data effectively, we can capture the temporal dynamics of patient movements. This adaptation involves incorporating time series-specific augmentations - such as time warping, jittering, and scaling - to enhance the model's ability to learn meaningful representations of sequential data. Additionally, employing neural network architectures suited for sequential data processing, such as one-dimensional convolutional neural networks (1D CNNs), ensures that the temporal structure of the data is preserved and effectively analysed.

Using these models enables the clustering of complex motion data collected from patients, allowing us to identify distinct movement patterns that may correspond to different neurological conditions or recovery trajectories. Traditional clinical evaluations, which rely on subjective assessments, can overlook subtle differences in motor performance that may be indicative of specific impairments. In contrast, machine learning models can process large amounts of pose data captured over extended periods, identifying patterns not easily observable by clinicians.

Furthermore, the integration of self-supervised learning reduces the need for manually labelled datasets, making it feasible to collect and analyse data at scale from stroke patients in their homes. As patients interact with their environments during daily activities, the algorithm can capture and analyse the subtleties of their movement in real time, without the need for constant clinician oversight. This level of granularity in phenotyping has the potential to reveal new categories of impairment and recovery, which in turn could guide more personalized rehabilitation programs. For example, identifying patients with similar movement patterns but different clinical outcomes could lead to insights about which rehabilitation strategies are most effective for particular phenotypes.

A.5 Summary of the Algorithm and Its Applications

The algorithm developed here leverages a combination of self-supervised learning (SimCLR), dimensionality reduction (UMAP), and density-based clustering (HDBSCAN) to analyse pose data captured from stroke patients during at-home evaluations. SimCLR allows for robust feature extraction from the raw pose data, while UMAP reduces the dimensionality of the learned features to facilitate clustering. HDBSCAN is used to group the data into meaningful clusters, identifying distinct motor patterns among patients.

The pipeline aims to be a scalable and automated method to post-stroke phenotyping, a complement to traditional, in-person clinical assessments. By capturing and analysing movement data in real-world environments, this algorithm aims to enrich the understanding of recovery processes with personalized rehabilitation strategies. The following sections provide a detailed breakdown of the algorithm, including the mathematical principles behind each component and all Python code.

Pipeline Overview

This section is offers a detailed explanation of the pipeline for pose data analysis, including the mathematical methods and corresponding Python code applied at each step. The process includes seven steps: Data Collection and Preprocessing, Feature Extraction, Self-Supervised Learning with SimCLR, Dimensionality

Reduction with UMAP, Clustering with HDBSCAN, Evaluation of Clustering, and visualisation of Cluster Health.

1. Data Collection and Preprocessing

In this step, we generate synthetic pose data that mimics realistic time series patterns, such as sine waves, cosine waves, sawtooth waves, and random noise. This data simulates human motion over time and is essential for training models that generalize well to real-world scenarios. Each sample consists of multiple time-steps and features, representing different pose coordinates over time.

Synthetic data generation involves creating time series $s(t)$ using functions like sine and cosine with added Gaussian noise to simulate measurement inaccuracies:

$$s(t) = f(t) + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2)$$

where $f(t)$ is a function like $\sin(t)$ or $\cos(t)$, and ε is Gaussian noise.

The generated data is then normalized per feature across all samples to ensure zero mean and unit variance, which is crucial for model training:

$$x' = \frac{x - \mu}{\sigma}$$

where x is the original data, μ is the mean, σ is the standard deviation, and x' is the normalized data.

```
import numpy as np
from sklearn.preprocessing import StandardScaler
from scipy import signal

# Generate synthetic pose data (time series data) with realistic patterns
def generate_synthetic_data(n_samples=1000, n_timesteps=120, n_features=17):
    np.random.seed(42)
    data = []
    for _ in range(n_samples):
```

```
        # Randomly choose a pattern
        pattern_type = np.random.choice(['sine', 'cosine', '
            sawtooth', 'random'])
        t = np.linspace(0, 2 * np.pi, n_timesteps)
        if pattern_type == 'sine':
            series = np.sin(t) + np.random.normal(0, 0.1, size
                =(n_timesteps,))
        elif pattern_type == 'cosine':
            series = np.cos(t) + np.random.normal(0, 0.1, size
                =(n_timesteps,))
        elif pattern_type == 'sawtooth':
            series = signal.sawtooth(t) + np.random.normal(0,
                0.1, size=(n_timesteps,))
        else:
            series = np.random.normal(0, 1, size=(n_timesteps,)
                )
        # Create features by adding slight variations
        sample = np.array([series + np.random.normal(0, 0.05,
            size=(n_timesteps,)) for _ in range(n_features)]).T
        data.append(sample)
    data = np.array(data)
    return data

data = generate_synthetic_data()

# Normalize the data per feature
n_samples, n_timesteps, n_features = data.shape
data_reshaped = data.reshape(-1, n_features)
scaler = StandardScaler()
data_normalized = scaler.fit_transform(data_reshaped)
data_normalized = data_normalized.reshape(n_samples,
    n_timesteps, n_features)
```

2. Feature Extraction

The normalized data is processed to extract key kinematic features such as position, velocity, acceleration, and jerk. These features capture the dynamics of motion and are essential for understanding the temporal patterns in the data. The derivatives are calculated without padding, accepting the reduced sequence length to avoid introducing artificial artifacts.

The derivatives are computed as:

$$v(t) = \frac{p(t+1) - p(t)}{\Delta t}$$

$$a(t) = \frac{v(t+1) - v(t)}{\Delta t}$$

$$j(t) = \frac{a(t+1) - a(t)}{\Delta t}$$

where: - $p(t)$ is the position at time t, - $v(t)$ is the velocity, - $a(t)$ is the acceleration, - $j(t)$ is the jerk, - Δt is the time step between observations (assumed to be 1 for simplicity).

The features are then concatenated along the feature axis to form a combined feature matrix.

```
def extract_features(data):
    # Compute velocity, acceleration, and jerk without
       padding
    velocity = np.diff(data, axis=1) # Shape: (n_samples,
       n_timesteps - 1, n_features)
    acceleration = np.diff(velocity, axis=1) # Shape: (
       n_samples, n_timesteps - 2, n_features)
    jerk = np.diff(acceleration, axis=1) # Shape: (n_samples,
        n_timesteps - 3, n_features)

    # Truncate data to match the shortest sequence
    min_timesteps = min(data.shape[1], velocity.shape[1],
       acceleration.shape[1], jerk.shape[1])
```

```
    data_truncated = data[:, :min_timesteps, :]
    velocity = velocity[:, :min_timesteps, :]
    acceleration = acceleration[:, :min_timesteps, :]
    jerk = jerk[:, :min_timesteps, :]

    # Concatenate features along the feature axis
    features = np.concatenate([data_truncated, velocity,
       acceleration, jerk], axis=-1)
    return features

features = extract_features(data_normalized)
n_samples, n_timesteps, n_features_combined = features.shape
```

3. Self-Supervised Learning with SimCLR

To learn robust representations without labelled data, we employ the SimCLR framework adapted for time series data. SimCLR uses contrastive learning by maximizing agreement between differently augmented views of the same data point. We use time series-specific augmentations such as time warping, jittering, and scaling.

The contrastive loss function used is the normalized temperature-scaled cross-entropy loss (NT-Xent):

$$\mathscr{L}_{i,j} = -\log \frac{\exp(\text{sim}(z_i, z_j)/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(\text{sim}(z_i, z_k)/\tau)}$$

where: - z_i, z_j are the normalized representations of augmented samples, - $\text{sim}(z_i, z_j) = z_i^\top z_j$ is the cosine similarity, - τ is the temperature parameter, - N is the batch size.

```
import tensorflow as tf
from tensorflow.keras import layers
from scipy import signal

# Time series-specific augmentations
def time_series_augmentation(batch):
```

```
    augmented_batch = []
    for series in batch:
        # Randomly apply augmentations
        if np.random.rand() < 0.5:
            # Time warping (stretching or squeezing)
            factor = np.random.uniform(0.8, 1.2)
            series = signal.resample(series, int(series.shape
                [0] * factor))
            # Pad or truncate to original length
            if series.shape[0] > n_timesteps:
                series = series[:n_timesteps, :]
            else:
                padding = n_timesteps - series.shape[0]
                series = np.pad(series, ((0, padding), (0, 0)),
                    'edge')
        if np.random.rand() < 0.5:
            # Jittering (adding noise)
            noise = np.random.normal(0, 0.05, series.shape)
            series = series + noise
        if np.random.rand() < 0.5:
            # Scaling
            scale = np.random.uniform(0.8, 1.2)
            series = series * scale
        augmented_batch.append(series)
    return np.array(augmented_batch, dtype=np.float32)

def simclr_loss(z_i, z_j, temperature=0.1):
    batch_size = tf.shape(z_i)[0]
    z = tf.concat([z_i, z_j], axis=0)
    z = tf.math.l2_normalize(z, axis=1)
    similarity_matrix = tf.matmul(z, z, transpose_b=True) /
        temperature
```

```
36
37     labels = tf.concat([tf.range(batch_size), tf.range(
           batch_size)], axis=0)
38     labels = tf.cast(labels, tf.int64)
39
40     mask = tf.eye(2 * batch_size, dtype=tf.bool)
41     logits = tf.boolean_mask(similarity_matrix, ~mask)
42     logits = tf.reshape(logits, [2 * batch_size, -1])
43
44     loss = tf.nn.sparse_softmax_cross_entropy_with_logits(
           labels=labels, logits=logits)
45     loss = tf.reduce_mean(loss)
46     return loss
47
48 class SimCLR(tf.keras.Model):
49     def __init__(self, encoder, projection_head):
50         super(SimCLR, self).__init__()
51         self.encoder = encoder
52         self.projection_head = projection_head
53
54     def call(self, x):
55         h = self.encoder(x)
56         z = self.projection_head(h)
57         return z
58
59 # Encoder adjusted for time series
60 encoder = tf.keras.Sequential([
61     layers.Conv1D(64, 3, activation='relu', input_shape=(
           n_timesteps, n_features_combined)),
62     layers.MaxPooling1D(),
63     layers.Conv1D(128, 3, activation='relu'),
64     layers.GlobalMaxPooling1D(),
```

```
    layers.Dense(256, activation='relu')
])

projection_head = tf.keras.Sequential([
    layers.Dense(128, activation='relu'),
    layers.Dense(128)
])

simclr_model = SimCLR(encoder, projection_head)
optimizer = tf.keras.optimizers.Adam()

# Training SimCLR on synthetic data
n_epochs = 10
batch_size = 32

for epoch in range(n_epochs):
    indices = np.arange(n_samples)
    np.random.shuffle(indices)
    for i in range(0, n_samples, batch_size):
        batch_indices = indices[i:i + batch_size]
        batch = features[batch_indices]
        with tf.GradientTape() as tape:
            augmented_i = time_series_augmentation(batch)
            augmented_j = time_series_augmentation(batch)
            z_i = simclr_model(augmented_i)
            z_j = simclr_model(augmented_j)
            loss = simclr_loss(z_i, z_j)
        gradients = tape.gradient(loss, simclr_model.
            trainable_variables)
        optimizer.apply_gradients(zip(gradients, simclr_model.
            trainable_variables))
    print(f"Epoch {epoch + 1}/{n_epochs}, Loss: {loss.numpy()
```

```
    }")
```

4. Dimensionality Reduction with UMAP

After obtaining the learned representations from the trained encoder, we apply Uniform Manifold Approximation and Projection (UMAP) to reduce the dimensionality of the data for clustering and visualisation. UMAP preserves the global and local structure of the data, making it suitable for high-dimensional embeddings.

```
import umap

# Obtain representations from the trained encoder
representations = encoder.predict(features)

# Apply UMAP to the representations
reducer = umap.UMAP(n_neighbors=15, min_dist=0.1, metric='
    euclidean')
embedding = reducer.fit_transform(representations)
```

5. Clustering with HDBSCAN

We perform clustering using the Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) algorithm. HDBSCAN identifies clusters of varying densities and sizes, and it can effectively handle noise in the data.

```
import hdbscan

# Clustering with HDBSCAN
clusterer = hdbscan.HDBSCAN(min_cluster_size=10)
clusters = clusterer.fit_predict(embedding)
```

6. Evaluation of Clustering

To evaluate the quality of the clusters, we use the Silhouette Score and the Davies-Bouldin Index (DBI). Since HDBSCAN labels noise points as -1, we exclude these points from the evaluation.

- **Silhouette Score**:

$$s(i) = \frac{b(i) - a(i)}{\max(a(i), b(i))}$$

where $a(i)$ is the mean intra-cluster distance for point i, and $b(i)$ is the mean nearest-cluster distance for point i.

- **Davies-Bouldin Index**:

$$\text{DBI} = \frac{1}{K} \sum_{i=1}^{K} \max_{j \neq i} \left(\frac{S_i + S_j}{M_{ij}} \right)$$

where K is the number of clusters, S_i is the average distance between each point in cluster i and the centroid of cluster i, and M_{ij} is the distance between the centroids of clusters i and j.

```
from sklearn.metrics import silhouette_score,
    davies_bouldin_score

# Filter out noise points for evaluation
labels = clusters[clusters >= 0]
embedding_filtered = embedding[clusters >= 0]

if len(np.unique(labels)) > 1:
    silhouette_avg = silhouette_score(embedding_filtered,
        labels)
    db_index = davies_bouldin_score(embedding_filtered,
        labels)
    stability_scores = clusterer.probabilities_[clusters >=
        0]
else:
    silhouette_avg = -1
    db_index = -1
    stability_scores = []

```

```
print("Silhouette Score:", silhouette_avg)
print("Davies-Bouldin Index:", db_index)
print("Cluster Stability Scores:", stability_scores)
```

7. visualisation of Cluster Health

Finally, we visualize the clustering results and assess the health of the clusters. We create UMAP projections of the clustered data and plot the cluster probabilities (stability scores) and sizes.

```
import matplotlib.pyplot as plt
import numpy as np

# Unique clusters and cluster sizes
unique_clusters = np.unique(clusters)
unique_clusters = unique_clusters[unique_clusters != -1]
cluster_sizes = np.bincount(clusters[clusters >= 0])

# UMAP Projection of HDBSCAN Clusters
plt.figure(figsize=(12, 8))
plt.scatter(embedding[:, 0], embedding[:, 1], c=clusters,
    cmap='Spectral', s=5, alpha=0.5)
plt.colorbar(boundaries=np.arange(len(set(clusters)) + 1) -
    0.5).set_ticks(np.arange(len(set(clusters))))
plt.title('UMAP Projection of HDBSCAN Clusters')
plt.xlabel('UMAP 1')
plt.ylabel('UMAP 2')
plt.show()

# HDBSCAN Cluster Stability
plt.figure(figsize=(10, 6))
plt.bar(unique_clusters, clusterer.probabilities_[clusters
    >= 0], color='skyblue')
plt.xlabel('Cluster Label')
```

```
plt.ylabel('Stability Score')
plt.title('HDBSCAN Cluster Stability')
plt.show()

# HDBSCAN Cluster Probability and Sizes
fig, ax1 = plt.subplots(figsize=(12, 6))
color = 'tab:blue'
ax1.set_xlabel('Cluster Label')
ax1.set_ylabel('Cluster Probability', color=color)
ax1.bar(unique_clusters, clusterer.probabilities_[clusters >= 0], color=color, alpha=0.6, label='Cluster Probability')
ax1.tick_params(axis='y', labelcolor=color)
ax1.legend(loc='upper left')

ax2 = ax1.twinx()
color = 'tab:green'
ax2.set_ylabel('Cluster Size', color=color)
ax2.plot(unique_clusters, cluster_sizes, color=color, marker='o', label='Cluster Size')
ax2.tick_params(axis='y', labelcolor=color)
ax2.legend(loc='upper right')

fig.tight_layout()
plt.title('HDBSCAN Cluster Probability and Sizes')
plt.show()
```

Notes on Synthetic Data

The integration of machine learning into healthcare presents a complex interplay between leveraging patient data for clinical advancement and the imperative to safeguard patient privacy. Synthetic data generation has emerged as a promising solution to this dilemma, offering the potential to utilize detailed datasets without

compromising sensitive information. However, the adoption of synthetic data introduces specific concerns that must be meticulously addressed to ensure ethical compliance and maintain the utility of such data in clinical research.

There exists an ethical imperative to harness all available technologies, including advanced machine learning algorithms, to enhance patient outcomes and healthcare delivery. Comprehensive datasets are essential for training robust machine learning models capable of uncovering novel insights and improving clinical decision-making processes. Synthetic data can enable researchers to exploit the statistical properties of real patient data while mitigating privacy risks, thus facilitating broader data sharing and collaboration. Despite these advantages, data custodians often perceive the risk-benefit balance as unfavourable due to legitimate privacy concerns, which restricts access to high-quality datasets necessary for advancing machine learning applications in healthcare.

The generation and application of synthetic data involve several complex challenges that must be addressed to realize their full potential. Defining data quality remains a significant issue, as there is a lack of consensus on metrics for assessing the fidelity of synthetic data to real datasets. For synthetic data to be valuable, they must accurately capture the joint distributions and complex relationships inherent in real data. Ensuring privacy preservation is also paramount; robust methods must be developed to prevent the inadvertent disclosure of patient information, satisfying both regulatory requirements and ethical standards. Furthermore, different clinical applications demand synthetic data with tailored characteristics, such as time-series data pertinent to stroke rehabilitation, which present unique challenges compared to static datasets used in other medical research areas.

In the context of the self-supervised learning pipeline for clinical phenotyping of pose data, the incorporation of synthetic data offers notable advantages. Synthetic data can supplement limited real-world datasets, enabling the training of more robust and generalizable machine learning models without the extensive data collection efforts often constrained by privacy concerns. By providing a diverse array of simulated patient movement patterns, synthetic data facilitate comprehensive

testing and validation of the algorithm's capability to capture and analyse complex motor impairments. Utilizing synthetic data aligns with ethical standards by reducing the risk of exposing sensitive patient information, thereby fostering greater trust among patients, clinicians, and data custodians.

Nevertheless, it is imperative to ensure that the generated synthetic pose data authentically reflect the nuanced motor impairments observed in post-stroke patients to be effective in refining phenotyping and tailoring rehabilitation strategies. Models trained on synthetic data should exhibit strong generalization when applied to real-world patient data to maintain clinical relevance and efficacy. Continuous evaluation and refinement of both the utility and privacy aspects of synthetic data are necessary to adapt to evolving challenges and to enhance data generation methodologies continually. By meticulously balancing the risks and benefits and directly addressing these technical and ethical challenges, the integration of synthetic data can significantly advance machine learning applications in healthcare, ultimately improving patient outcomes and contributing to a more effective healthcare system.

Conclusion

This pipeline illustrates the processing and analysis of pose data using time series-specific methods. By generating synthetic data and leveraging feature extraction and augmentation techniques, this approach integrates an adapted SimCLR framework for time series data with UMAP for dimensionality reduction and HDBSCAN for clustering to identify patterns in pose data. Evaluation metrics and visualisations further assess the quality and stability of the clustered behaviours.

Appendix B

HIERARCHICAL CLUSTERING WITH GOWER DISTANCE FOR DYNAMIC PATIENT GROUPING

B.1 Rationale

Stroke is a leading cause of long-term disability worldwide, significantly impacting survivors' motor, sensory, and cognitive functions (Feigin et al., 2021). Rehabilitation aims to restore these functions, but outcomes vary widely among individuals. One critical yet under-explored factor influencing rehabilitation efficacy is the social dynamics within the rehabilitation setting. Research suggests that patients may benefit from interacting with peers who share similar levels of impairment, as this can enhance motivation, foster mutual understanding, and allow for appropriately matched therapeutic activities (Bandura, 1977; English and Hillier, 2011). Conversely, training with individuals of significantly different impairment levels may present challenges, such as feelings of frustration or decreased self-efficacy (Nichols-Larsen et al., 2005).

Traditional rehabilitation programs often overlook the social and psychological dimensions of recovery, focusing primarily on physical therapy. However, the composition of rehabilitation groups can significantly affect patient motivation, engagement, and overall progress (Salbach et al., 2006). Patients training alongside

peers with similar impairment levels tend to show improved motivation and engagement, leading to better functional outcomes. In contrast, heterogeneous groups with wide variations in impairment levels may hinder progress for some patients due to mismatched therapeutic activities and social comparison effects (Törnbom et al., 2019). Therefore, understanding and curating patients' psychosocial support structures is essential for optimizing rehabilitation outcomes. Personalizing group rehabilitation by considering patients' impairment levels, cognitive abilities, and social capacities can enhance rehabilitation efficacy.

B.2 Machine Learning in Patient Grouping

Advancements in machine learning (ML) and data analytics offer new possibilities for personalising rehabilitation groupings. ML algorithms can analyse complex patient data, including impairment severity, cognitive function, and social factors, to identify patterns and clusters among patients (Xu and Wunsch, 2005; Chen et al., 2015). Clustering algorithms, such as k-means clustering or hierarchical clustering, can group patients based on multiple variables, facilitating the creation of rehabilitation groups with similar needs. Applying ML to patient grouping involves several challenges, including data collection, feature selection, and algorithm choice. Collecting complete patient data requires integrating clinical assessments, psychological evaluations, and possibly in-situ time series. Feature selection must balance the inclusion of relevant variables with the complexity of the model to avoid overfitting.

Unsupervised learning approaches like clustering are suitable for this task, as they do not require labelled data and can reveal natural groupings within the patient population (Jain, 2010). By clustering patients based on impairment levels, cognitive function scores, and social engagement metrics, clinicians can form groups that maximize therapeutic effectiveness and patient satisfaction. Additionally, ML models can incorporate psychological assessments to evaluate patients' social capacities and preferences, further refining group assignments. By leveraging these technologies, clinicians can design more effective, personalized rehabilitation programs that account for the multifaceted nature of patient recovery.

B.3 Key Considerations

Implementing a machine learning algorithm in a hospital setting to dynamically calculate the fit between patients involves several key considerations. Firstly, the characteristics of the data must be carefully evaluated. Patient data typically includes mixed data types, such as numerical values (e.g., impairment scores), categorical variables (e.g., types of impairments), and possibly textual data from clinical notes. The data may also exhibit high dimensionality due to the numerous factors involved, such as physical impairment levels, cognitive abilities, psychological assessments, and social factors. Additionally, missing data, or "missingness", is a common issue in clinical settings, which must be addressed to ensure the algorithm's effectiveness.

Secondly, clinical requirements play a crucial role in selecting and implementing the algorithm. Interpretability is essential; clinicians need to understand how the algorithm reaches its conclusions to trust and effectively use its recommendations. Scalability and efficiency are important because the algorithm must handle data from many patients to be practical in a busy hospital environment. Privacy and compliance with healthcare regulations (e.g., HIPAA, GDPR) are mandatory, necessitating secure handling of patient data. Furthermore, the algorithm should support dynamic updating to adapt as new patient data becomes available, ensuring that patient groupings remain relevant over time.

B.4 Suitable Machine Learning Algorithms

Several machine learning algorithms are suitable for grouping patients based on key factors reported in the literature. Unsupervised clustering algorithms are particularly appropriate, as they can identify natural groupings within the data without requiring labelled outcomes. Among these, hierarchical clustering with Gower distance, K-Prototypes clustering, fuzzy clustering, and graph-based clustering are notable candidates for this task.

B.4.1 Hierarchical Clustering with Gower Distance

Hierarchical clustering builds a tree-like structure (dendrogram) of clusters without needing to specify the number of clusters in advance. The Gower distance metric

measures similarity between records with mixed data types, handling both numerical and categorical variables. This method provides interpretable dendrograms, allowing clinicians to visualise and understand patient groupings. However, it can be computationally intensive with very large datasets.

B.4.2 K-Prototypes Clustering

K-Prototypes clustering is an extension of K-Means clustering designed to handle mixed data types by combining K-Means (for numerical data) and K-Modes (for categorical data) algorithms. The algorithm requires specifying the number of clusters (K) and may converge to local minima, necessitating multiple runs. It is suitable when the approximate number of patient groups is known from clinical experience.

B.4.3 Fuzzy Clustering

Fuzzy clustering algorithms, such as Fuzzy C-Means, assign membership levels to data points for multiple clusters, reflecting the degree to which each patient belongs to each group. This approach acknowledges that patients may share characteristics with multiple groups, providing a nuanced view of patient similarities. It requires specifying the number of clusters and can be more complex to interpret compared to hard clustering methods. Fuzzy clustering is useful when patient characteristics overlap, and flexible group assignments are beneficial.

B.4.4 Graph-Based Clustering

Graph-based clustering models patients as nodes in a graph, with edges representing similarities between them. Clusters are identified as communities within the graph using community detection algorithms. This method captures complex relationships and is flexible in handling different data types. Graph-based clustering is particularly suitable when relationships between patients are complex and not easily captured by traditional distance measures. It is effective in situations where patient similarities involve intricate patterns or when incorporating relational data, such as social interactions or co-morbidities. This method can reveal community structures within patient populations that may inform effective groupings.

B.5 Summary of the Algorithm and Its Applications

This section presents a detailed explanation of the algorithm for patient grouping using hierarchical clustering with Gower distance. The algorithm is designed to handle mixed data types (numerical and categorical) commonly found in patient data. The pipeline includes several steps: data collection and preprocessing, feature selection and encoding, calculation of the Gower distance matrix, hierarchical clustering, determination of optimal clusters, evaluation and visualisation of clusters, and dynamic updating. The proposed algorithm aims to identify meaningful patient groupings based on multifaceted clinical factors.

Pipeline Overview

The pipeline begins with the collection and preprocessing of patient data, ensuring that all relevant numerical and categorical variables are accurately captured and cleaned. Following this, feature selection and encoding are performed in collaboration with clinicians to identify and prepare the most significant factors influencing rehabilitation outcomes. Feature engineering and dimensionality reduction techniques are applied to optimize the dataset for effective clustering. The Gower distance matrix is then calculated to measure similarities between patients with mixed data types. Hierarchical clustering is applied to this distance matrix to form initial groupings, after which the optimal number of clusters is determined using validation metrics. The clusters are then evaluated and visualized to assess their cohesion, separation, and clinical relevance. Finally, dynamic updating methods are incorporated to ensure that the clustering model remains current as new patient data become available and existing patients progress through rehabilitation.

1. Data Collection and Preprocessing

In this step, patient data are collected and prepared for analysis. Patient data are sourced from medical records and recent assessments, ensuring comprehensive coverage of all relevant features. The dataset includes both numerical and categorical variables relevant to rehabilitation grouping, such as impairment levels measured by standardized scales like the Fugl-Meyer Assessment, cognitive function assessed

via tests such as the Mini-Mental State Examination (MMSE) scores, social factors including support systems ('Weak', 'Moderate', 'Strong'), communication abilities ('Poor', 'Fair', 'Good'), and psychological assessments covering motivation levels ('Low', 'Medium', 'High'), self-efficacy, and mood evaluations.

For numerical features such as impairment level and cognitive scores, missing values are imputed using the mean or median of the available data to maintain statistical consistency. For categorical features like support system strength, communication ability, and motivation level, missing values are imputed using the mode (the most frequent category) or by introducing a new category to represent missing data, thereby preserving the integrity of categorical distributions.

```
import pandas as pd
import numpy as np

# Load the dataset
data = pd.read_csv('patient_data.csv')

# Handle missing values
# For numerical features
numerical_features = ['Impairment_Level', 'Cognitive_Score']
data[numerical_features] = data[numerical_features].fillna(
    data[numerical_features].mean())

# For categorical features
categorical_features = ['Support_System', '
    Communication_Ability', 'Motivation_Level']
data[categorical_features] = data[categorical_features].
    fillna(data[categorical_features].mode().iloc[0])
```

2. Feature Selection and Encoding

Features relevant to the grouping are selected, and categorical variables are encoded into numerical representations suitable for distance calculations.

Feature Selection

Building upon the prepared data, relevant features are selected in collaboration with clinicians to ensure that the most significant factors influencing rehabilitation outcomes are included. This process involves assessing the relevance of each feature and employing feature selection techniques, such as Recursive Feature Elimination (RFE) or feature importance scores from tree-based models, to systematically evaluate and select pertinent features. This process ensures that the clustering algorithm focuses on variables that have a meaningful impact on rehabilitation outcomes, thereby enhancing the quality and interpretability of the resulting clusters.

Feature Encoding

Categorical variables need to be converted into numerical formats suitable for processing by the clustering algorithm. Label encoding is employed to assign a unique integer to each category, facilitating the inclusion of categorical data in distance calculations. Specifically, support system strength is encoded as 'Weak' $\rightarrow$ 0, 'Moderate' $\rightarrow$ 1, and 'Strong' $\rightarrow$ 2. Communication ability follows a similar encoding scheme: 'Poor' $\rightarrow$ 0, 'Fair' $\rightarrow$ 1, and 'Good' $\rightarrow$ 2. Motivation level is encoded as 'Low' $\rightarrow$ 0, 'Medium' $\rightarrow$ 1, and 'High' $\rightarrow$ 2. This encoding process transforms categorical data into a numerical format that retains the ordinal nature of the categories, preserving meaningful relationships between them.

```
from sklearn.preprocessing import LabelEncoder

# Initialize LabelEncoders
le_support = LabelEncoder()
le_communication = LabelEncoder()
le_motivation = LabelEncoder()

# Encode categorical features
data['Support_System_Encoded'] = le_support.fit_transform(
    data['Support_System'])
data['Communication_Encoded'] = le_communication.
```

```
    fit_transform(data['Communication_Ability'])
data['Motivation_Encoded'] = le_motivation.fit_transform(
    data['Motivation_Level'])
```

Feature Engineering

Beyond selecting existing features, feature engineering can create new variables that capture underlying patterns or interactions between existing features. For instance, combining cognitive function scores with motivation levels may yield a composite feature that better represents a patient's readiness for rehabilitation. Such engineered features can provide additional insights and improve the clustering performance by highlighting more nuanced patient characteristics.

```
# Example of feature engineering by creating interaction
    terms
data['Cognitive_Motivation'] = data['Cognitive_Score'] *
    data['Motivation_Encoded']
```

Dimensionality Reduction

If the number of features is large, dimensionality reduction techniques may be necessary to enhance the efficiency and performance of the clustering algorithm. Principal Component Analysis (PCA) is one such technique that transforms the original features into a set of linearly uncorrelated components while preserving as much variance as possible. By reducing dimensionality, PCA helps mitigate the curse of dimensionality, making the clustering process more manageable and interpretable.

PCA seeks to find the principal components $\mathbf{PC}_1, \mathbf{PC}_2, \ldots, \mathbf{PC}_k$ that maximize the variance in the data:

$$\mathbf{PC}_i = \mathbf{w}_i^\top \mathbf{X}$$

where $\mathbf{w}_i$ are the eigenvectors corresponding to the largest eigenvalues of the covariance matrix of $\mathbf{X}$, the feature matrix.

```
from sklearn.decomposition import PCA
```

```
2
3 # Initialize PCA to retain 95% of the variance
4 pca = PCA(n_components=0.95, random_state=42)
5 X_reduced = pca.fit_transform(data[features])
6
7 print(f"Original number of features: {data[features].shape
    [1]}")
8 print(f"Reduced number of features after PCA: {X_reduced.
    shape[1]}")
```

3. Calculation of Gower Distance Matrix

Gower distance is used to measure similarity between records containing mixed data types. The Gower distance between two samples i and j is calculated as:

$$D_{ij} = \frac{\sum_{k=1}^{p} w_k d_{ijk}}{\sum_{k=1}^{p} w_k}$$

where:

- p is the number of features.
- w_k is the weight of feature k (usually 1).
- d_{ijk} is the distance between samples i and j for feature k:
  - For numerical features:

$$d_{ijk} = \frac{|x_{ik} - x_{jk}|}{\text{range}_k}$$

  - For categorical features:

$$d_{ijk} = \begin{cases} 0 & \text{if } x_{ik} = x_{jk} \\ 1 & \text{if } x_{ik} \neq x_{jk} \end{cases}$$

```
import gower # Make sure to install the 'gower' library

# Prepare data for Gower distance
features = ['Impairment_Level', 'Cognitive_Score', '
    Support_System_Encoded', 'Communication_Encoded', '
    Motivation_Encoded', 'Cognitive_Motivation']
X = data[features]

# Calculate Gower distance matrix
gower_dist = gower.gower_matrix(X)
```

4. Hierarchical Clustering

Based on the data characteristics and clinical requirements, hierarchical clustering with Gower distance is selected. This method is suitable for handling mixed data types without a predefined number of clusters.

```
from scipy.cluster.hierarchy import linkage, dendrogram

# Perform hierarchical clustering
Z = linkage(gower_dist, method='ward')

# Plot dendrogram
import matplotlib.pyplot as plt

plt.figure(figsize=(10, 7))
dendrogram(Z)
plt.title('Hierarchical Clustering Dendrogram')
plt.xlabel('Sample Index')
plt.ylabel('Distance')
plt.show()
```

5. Determination of Optimal Clusters

The optimal number of clusters k is determined by analysing the dendrogram or using methods like the Elbow method or silhouette scores.

```
from scipy.cluster.hierarchy import fcluster

# Decide on the number of clusters
k = 3 # Example number of clusters; in practice, determine based on analysis

# Assign cluster labels
cluster_labels = fcluster(Z, k, criterion='maxclust')
data['Cluster'] = cluster_labels
```

6. Evaluation and visualisation of Clusters

Clustering results are first evaluated using internal validation metrics, such as silhouette scores, to assess cluster cohesion and separation. Second, clinicians evaluate the clustering results to ensure interpretability and can thereafter guide adjustments to the model or data preprocessing steps.

Silhouette Score

The silhouette score $s(i)$ for each sample i is calculated as:

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}$$

where:

- $a(i)$ is the mean intra-cluster distance for sample i.
- $b(i)$ is the mean nearest-cluster distance for sample i.

```
from sklearn.metrics import silhouette_score

# Since we have a precomputed distance matrix, use metric='precomputed'
```

```
silhouette_avg = silhouette_score(gower_dist, cluster_labels
    , metric='precomputed')
print("Silhouette Score:", silhouette_avg)
```

Visualisation

Multidimensional Scaling (MDS) is used to reduce the dimensionality of the distance matrix for plotting, facilitating the visualisation of clusters.

Note: The visualisation step uses MDS to project high-dimensional data into two dimensions, allowing for graphical representation of the clusters.

```
from sklearn.manifold import MDS

# Reduce dimensions for visualisation
mds = MDS(n_components=2, dissimilarity='precomputed',
    random_state=42)
coords = mds.fit_transform(gower_dist)

# Plot clusters
plt.figure(figsize=(10, 7))
scatter = plt.scatter(coords[:, 0], coords[:, 1], c=
    cluster_labels, cmap='Set1')
plt.legend(*scatter.legend_elements(), title='Clusters')
plt.title('Hierarchical Clustering with Gower Distance')
plt.xlabel('MDS Dimension 1')
plt.ylabel('MDS Dimension 2')
plt.show()
```

7. Dynamic Updating

To ensure that patient groupings remain current and relevant, the clustering model must be capable of dynamically adjusting as new patient data become available and as existing patients progress through rehabilitation. Several methods and algorithms can facilitate this dynamic updating, including online learning techniques,

incremental clustering algorithms, and the integration of clinician feedback.

Online learning algorithms are designed to update models incrementally as new data points arrive, without the need to retrain the entire model from scratch. For clustering, incremental algorithms such as Incremental Hierarchical Clustering, Online K-Means, or density-based methods with incremental capabilities can be employed. These methods adjust clusters to accommodate new data points or changes in existing data. In the context of hierarchical clustering with Gower distance, incremental hierarchical clustering can be adapted to handle new patient data. As new patients are admitted or existing patients' assessments are updated, the algorithm recalculates distances involving the new or updated data points and adjusts the dendrogram accordingly. Given an existing distance matrix D and a new data point x_{new}, the updated distance matrix D' can be computed by calculating the Gower distances between x_{new} and all existing data points:

$$D' = \begin{bmatrix} D & d_{\text{new}} \\ d_{\text{new}}^{\top} & 0 \end{bmatrix}$$

where d_{new} is the vector of distances between x_{new} and each existing data point.

Bayesian clustering methods offer a probabilistic framework that incorporates prior information into the clustering process. By placing prior distributions over model parameters - such as cluster means, variances, or assignments - the approach allows integration of previous patient scores or clinician insights as prior beliefs. These priors are then updated with observed data to form posterior distributions, resulting in clusters that reflect both existing knowledge and new evidence.

The posterior probability $P(\theta|X)$ is proportional to the product of the likelihood $P(X|\theta)$ and the prior $P(\theta)$:

$$P(\theta|X) \propto P(X|\theta) \cdot P(\theta)$$

where θ represents the cluster parameters, and X is the data.

Semi-supervised learning methods allow the incorporation of expert knowledge by using a combination of labelled and unlabelled data. Constraint-based

clustering algorithms modify the clustering process to satisfy clinician-provided constraints. The constrained clustering problem can be formulated by introducing a penalty term in the clustering objective function that accounts for the constraints:

$$\text{Objective} = \text{Clustering Loss} + \lambda \times \text{Constraint Penalty}$$

where λ controls the influence of the constraints.

While many clustering algorithms are not inherently incremental, libraries such as `scikit-learn` provide implementations for incremental versions of algorithms like K-Means. For hierarchical clustering with mixed data types, custom implementations may be required.

```
# Updating the Gower distance matrix with new data
import numpy as np
import gower

# Existing data
X_existing = data[features]
gower_dist_existing = gower.gower_matrix(X_existing)

# New data point
X_new = new_data[features]

# Calculate distances between new data and existing data
gower_dist_new = gower.gower_matrix(np.vstack([X_existing,
    X_new]))[-1, :-1]

# Update the distance matrix
gower_dist_updated = np.vstack([
    np.hstack([gower_dist_existing, gower_dist_new.reshape
        (-1, 1)]),
    np.hstack([gower_dist_new, np.array([0])])
])
```

```
20
21 # Perform hierarchical clustering on the updated distance
      matrix
22 from scipy.cluster.hierarchy import linkage
23
24 Z_updated = linkage(gower_dist_updated, method='ward')
```

Selecting the appropriate method for dynamic updating depends on factors such as the frequency of data updates, dataset size, computational resources, and the importance of incorporating prior information and clinician feedback. A hybrid approach may be most effective, combining incremental clustering for handling new data with periodic retraining to recalibrate the model.

A Note on Dynamic Updating

While hierarchical clustering does not natively support incremental updates, updating the distance matrix and recomputing the clustering allows for the inclusion of new data. For large datasets, recomputing the clustering can be computationally intensive; therefore, considering alternative clustering algorithms that support incremental updates may be necessary if scalability becomes an issue. Incorporating clinician feedback through semi-supervised learning enhances the clinical relevance of the clusters. Additionally, saving model components and data facilitates seamless updates as new information becomes available, ensuring the model remains current.

Conclusion

This pipeline highlights the application of hierarchical clustering with Gower distance to uncover meaningful patient groupings based on multifaceted clinical factors. Through careful data preparation, feature selection, and variable engineering, the dataset is optimised to enhance clustering performance. The inclusion of dynamic updating ensures that the model remains adaptable, continuously refining clusters in response to new information about patient status. This approach offers a flexible and insightful framework for analysing complex clinical datasets.

References

Abdelghani, M., Lillicrap, T., and Tweed, D. (2008). Sensitivity derivatives for flexible sensorimotor learning. *Neural Comput.*, 20:2085–2111.

Abeele, S. and Bock, O. (2001). Sensorimotor adaptation to rotated visual input: Different mechanisms for small versus large rotations. *Exp. Brain Res.*, 140:407–410.

Abraham, A. and Drory, V. E. (2014). Listening to music during electromyography does not influence the examinee's anxiety and pain levels. *Muscle & Nerve*, 50(3):445–447.

Abrahamse, E., Jiménez, L., Verwey, W., and Clegg, B. (2010). Representing serial action and perception. *Psychon. Bull. Rev.*, 17:603–623.

Abrahamse, E., Ruitenberg, M., de Kleine, E., and Verwey, W. (2013). Control of automated behavior: Insights from the discrete sequence production task. *Front. Hum. Neurosci.*, 7:82.

Adams, J. A. (1971). A closed-loop theory of motor learning. *Journal of Motor Behavior*, 3:111–150.

Adams, J. A. (1987). Historical review and appraisal of research on the learning, retention, and transfer of human motor skills. *Psychological Bulletin*, 101(1):41–74.

Aeles, J., Kelly, L. A., Yoshitake, Y., and Cresswell, A. G. (2020). Fine-wire recordings of flexor hallucis brevis motor units up to maximal voluntary contraction

reveal a flexible, non-rigid mechanism for force control. *Journal of Neurophysiology*, 123:1766–1774.

Ahmadi-Pajouh, M., Towhidkhah, F., and Shadmehr, R. (2012). Preparing to reach: Selecting an adaptive long-latency feedback controller. *J. Neurosci.*, 32:9537–9545.

Albouy, G., Sterpenich, V., Balteau, E., et al. (2008). Both the hippocampus and striatum are involved in consolidation of motor sequence memory. *Neuron*, 58:261–272.

Albouy, G., Sterpenich, V., Vandewalle, G., et al. (2013). Interaction between hippocampal and striatal systems predicts subsequent consolidation of motor sequence memory. *PLoS One*, 8:e59490.

Ambike, S., Huang, X., and Latash, M. L. (2021). Synergic control in neurological patients: Insights from studies of motor impairments. *Motor Control*, 25:442–457.

Antonenko, D., Diekelmann, S., Olsen, C., Born, J., and Mölle, M. (2013). Napping to renew learning capacity: enhanced encoding after stimulation of sleep slow oscillations. *European Journal of Neuroscience*, 37(7):1142–1151.

Arac, A., Zhao, P., Dobkin, B. H., Carmichael, S. T., and Golshani, P. (2019). Deepbehavior: A deep learning toolbox for automated analysis of animal and human behavior imaging data. *Frontiers in Systems Neuroscience*, 13:20.

Ariani, G. and Diedrichsen, J. (2019). Sequence learning is driven by improvements in motor planning. *Journal of Neurophysiology*, 122(1):135–147.

Babinski, J. (1899). Sur le signe du pied et sur les reflexes cutanes. *Archives de Neurologie*, 9:317–331.

Baddeley, A. D. and Longman, D. J. (1978). The influence of length and spacing of training sessions on the rate of learning to type. *Ergonomics*, 21:627–635.

Balleine, B. W. and Dickinson, A. (1998). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37:407–419.

Bandura, A. (1977). *Social Learning Theory*. Prentice-Hall, Englewood Cliffs, NJ.

Barakat, M., Doyon, J., Debas, K., and et al. (2013). Sleep spindles predict neural and behavioral changes in motor sequence consolidation. *Human Brain Mapping*, 34(11):2918–2928.

Bellman, R. E. (1957). *Dynamic Programming*. Princeton University Press, Princeton, NJ.

Bellman, R. E. (1961). *Adaptive Control Processes: A Guided Tour*. Princeton University Press, Princeton, NJ.

Benke, T. (1993). Two forms of apraxia in alzheimer's disease. *Cortex*, 29:715–725.

Berg, R., Caggiano, V., and Kumar, V. (2023). Sar: Generalization of physiological agility and dexterity via synergistic action representation. *arXiv preprint arXiv:2307.03716*.

Berman, G. J., Choi, D. M., Bialek, W., and Shaevitz, J. W. (2014). Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface*, 11(99):20140672.

Berniker, M. and Kording, K. P. (2008). Estimating the sources of motor errors for adaptation and generalization. *Nature Neuroscience*, 11:1454–1461.

Berniker, M. and Kording, K. P. (2011). Estimating the relevance of world disturbances to explain savings, interference and long-term motor adaptation effects. *PLOS Computational Biology*, 7(10):e1002210.

Bernstein, N. (1947a). *On Dexterity and Its Development*. Lawrence Erlbaum Associates.

Bernstein, N. (1947b). *On the Construction of Movements*. Medgiz, Moscow, Russia.

Bernstein, N. A. (1967). *The co-ordination and regulation of movements*. Pergamon Press.

Bhattacharyya, A. (1946). On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society*, 35:99–109.

Bird, G. and Heyes, C. M. (2005). Effector-dependent learning by observation of a finger movement sequence. *Journal of Experimental Psychology: Human Perception and Performance*, 31(2):262–275.

Bizzi, E., Cheung, V. C., d'Avella, A., Saltiel, P., and Tresch, M. (2008). Combining modules for movement. *Brain Research Reviews*, 57(1):125–133.

Bizzi, E. and Cheung, V. C. K. (2013). The neural origin of muscle synergies. *Frontiers in Computational Neuroscience*, 7:51.

Bizzi, E., Mussa-Ivaldi, F. A., and Giszter, S. F. (1991). Computations underlying the execution of movement: A biological perspective. *Science*, 253(5017):287–291.

Bjerknes, S., Toft, M., Brandt, R., Rygvold, T. W., Konglund, A., Dietrichs, E., Andersson, S., and Skogseid, I. M. (2021). Subthalamic nucleus stimulation in parkinson's disease: 5-year extension study of a randomized trial. *Movement Disorders Clinical Practice*, 9(1):48–59.

Bobath, B. (1978). *Adult Hemiplegia: Evaluation and Treatment*. London: Heinemann Medical.

Bonstrup, M., Iturrate, I., Thompson, R., Cruciani, G., Censor, N., and Cohen, L. G. (2019). A rapid form of offline consolidation in skill learning. *Current Biology*, 29(8):1346–1351.e4.

Born, J. and Wilhelm, I. (2010). System consolidation of memory during sleep. *Psychological Research*, 24:122–137.

Boyd, L. A., Edwards, J. D., Siengsukon, C. S., Vidoni, E. D., Wessel, B. D., and Linsdell, M. A. (2009). Motor sequence chunking is impaired by basal ganglia stroke. *Neurobiology of Learning and Memory*, 92(1):35–44.

Brunton, L. L., Chabner, B. A., and Knollmann, B. C. (2009). Goodman and gilman's: The pharmacological basis of therapeutics. *McGraw-Hill Medical*.

Bruton, M. and O'Dwyer, N. (2018). Synergies in coordination: a comprehensive overview of neural, computational, and behavioral approaches. *Journal of Neurophysiology*, 120(6):2761–2774.

Buch, E. R., Claudino, L., Quentin, R., Bönstrup, M., and Cohen, L. G. (2021). Consolidation of human skill linked to waking hippocampo-neocortical replay. *Cell Reports*, 35(10):109193.

Campello, R. J. G. B., Moulavi, D., and Sander, J. (2013). Density-based clustering based on hierarchical density estimates. In *Advances in Knowledge Discovery and Data Mining. PAKDD 2013*, volume 7819 of *Lecture Notes in Computer Science*. Springer, Berlin, Heidelberg.

Cao, Z., Hidalgo, G., Simon, T., Wei, S.-E., and Sheikh, Y. (2021). Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):172–186.

Cappellini, G., Ivanenko, Y. P., Dominici, N., Poppele, R. E., and Lacquaniti, F. (2010). Migration of motor pool activity in the spinal cord reflects body mechanics in human locomotion. *Journal of Neurophysiology*, 104(6):3064–3073.

Castellini, C. and van der Smagt, P. (2013). Evidence of muscle synergies during human grasping. *Biological Cybernetics*, 107(2):233–245.

Censor, N., Dayan, E., and Cohen, L. G. (2014). Cortico-subcortical neuronal circuitry associated with reconsolidation of human procedural memories. *Cortex*, 58:281–288.

Chase, W. G. and Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4:55–81.

Chen, J., Jin, W., Zhang, X., Xu, W., Liu, X., and Ren, C. (2015). Telerehabilitation approaches for stroke patients: Systematic review and meta-analysis of randomized controlled trials. *Journal of Stroke and Cerebrovascular Diseases*, 24(12):2660–2668.

Cheng, B., Xiao, B., Wang, J., Shi, H., Huang, T. S., and Zhang, L. (2020). Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation. *arXiv preprint arXiv:1908.10357*. CVPR 2020.

Cheung, V. C. K., d'Avella, A., Tresch, M. C., and Bizzi, E. (2005). Central and sensory contributions to the activation and organization of muscle synergies during natural motor behaviors. *Journal of Neuroscience*, 25(27):6419–6434.

Cheung, V. C. K., Piron, L., Agostini, M., Silvoni, S., Turolla, A., and Bizzi, E. (2009). Stability of muscle synergies for voluntary actions after cortical stroke in humans. *Proceedings of the National Academy of Sciences*, 106(46):19563–19568.

Christiansen, L., Madsen, M. J., Bojsen-Møller, E., Thomas, R., Nielsen, J. B., and Lundbye-Jensen, J. (2018). Progressive practice promotes motor learning and repeated transient increases in corticospinal excitability across multiple days. *Brain Stimulation*, 11(2):346–357.

Cirstea, C. M. and Levin, M. F. (2000). Compensatory strategies for reaching in stroke. *Brain*, 123(5):940–953.

Classen, J., Liepert, J., Wise, S. P., Hallett, M., and Cohen, L. G. (1998). Rapid plas-

ticity of human cortical movement representation induced by practice. *Journal of Neurophysiology*, 79(2):1117–1123.

Cossu, G. and Pau, M. (2017). Subthalamic nucleus stimulation and gait in parkinson's disease: A not always fruitful relationship. *Gait & Posture*, 52:205–210.

Danion, F., Latash, M. L., Li, Z. M., and Zatsiorsky, V. M. (2000). The effect of fatigue on multifinger co-ordination in force production tasks in humans. *The Journal of Physiology*, 523(Pt 2):523–532.

d'Avella, A. and Bizzi, E. (2005). Shared and specific muscle synergies in natural motor behaviors. *Proceedings of the National Academy of Sciences of the United States of America*, 102(8):3076–3081.

d'Avella, A., Portone, A., Fernandez, L., and Lacquaniti, F. (2006). Control of fast-reaching movements by muscle synergy combinations. *Journal of Neuroscience*, 26(30):7791–7810.

Daw, N. D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8:1704–1711.

Dayan, E. and Cohen, L. G. (2011). Neuroplasticity subserving motor skill learning. *Neuron*, 72(3):443–454.

De Freitas, P. B., Freitas, S. M. S. F., Reschechtko, S., Corson, T., Lewis, M. M., Huang, X., and Latash, M. L. (2020a). Synergic control of action in levodopa-naïve parkinson's disease patients: I. multi-finger interaction and coordination. *Experimental Brain Research*, 238(2):229–245.

De Freitas, P. B., Freitas, S. M. S. F., Reschechtko, S., Corson, T., Lewis, M. M., Huang, X., and Latash, M. L. (2020b). Synergic control of action in levodopa-naïve parkinson's disease patients: Ii. multi-muscle synergies stabilizing vertical posture. *Experimental Brain Research*, 238(2):2931–2945.

Delis, I., Panzeri, S., Pozzo, T., and Berret, B. (2018). Muscle synergy analysis of a hand-grasp dataset. *Frontiers in Neurorobotics*, 12:57.

Dewald, J. P., Pope, P. S., Given, J. D., Buchanan, T. S., and Rymer, W. Z. (1995). Abnormal muscle coactivation patterns during isometric torque generation at the elbow and shoulder in hemiparetic subjects. *Brain*, 118(Pt 2):495–510.

Dickinson, A. and Balleine, B. W. (1994). Motivational control of goal-directed action. *Animal Learning and Behavior*, 43:41–51.

Diedrichsen, J., Hashambhoy, Y., Rane, T., and Shadmehr, R. (2005). Neural correlates of reach errors. *Journal of Neuroscience*, 25(43):9919–9931.

Diedrichsen, J., King, M., Hernandez-Castillo, C., Sereno, M. I., and Ivry, R. B. (2019). Universal transform or multiple functionality? understanding the contribution of the human cerebellum across task domains. *Neuron*, 102(5):918–928.

Diedrichsen, J. and Kornysheva, K. (2015). Motor skill learning between selection and execution. *Trends in Cognitive Sciences*, 19(4):227–233.

Diekelmann, S. and Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*, 11(2):114–126.

Donchin, O., Francis, J. T., and Shadmehr, R. (2003). Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: Theory and experiments in human motor control. *Journal of Neuroscience*, 23(27):9032–9045.

Dong, L., Shen, B., Pang, Y., Zhang, M., Xiang, Y., Xing, Y., Wright, M., Li, D., and Bo, J. (2021). Fms effects of a motor program for children with autism spectrum disorders. *Perceptual and Motor Skills*, 128(4):1421–1442.

Doyon, J. and Benali, H. (2005). Reorganization and plasticity in the adult brain during learning of motor skills. *Current Opinion in Neurobiology*, 15(2):161–167.

Doyon, J., Penhune, V., and Ungerleider, L. G. (2003). Distinct contribution of the cortico-striatal and cortico-cerebellar systems to motor skill learning. *Neuropsychologia*, 41(3):252–262.

Doyon, J. and Ungerleider, L. G. (2002). Functional anatomy of motor skill learning. In Squire, L. R. and Schacter, D. L., editors, *Neuropsychology of Memory*, pages 225–238. Guilford Press.

English, C. and Hillier, S. (2011). Circuit class therapy for improving mobility after stroke: a systematic review. *Journal of Rehabilitation Medicine*, 43(7):565–571.

Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96)*, pages 226–231. AAAI Press.

Falaki, A., Cuadra, C., Lewis, M. M., Prado-Rico, J. M., Huang, X., and Latash, M. L. (2023). Multi-muscle synergies in preparation for gait initiation in parkinson's disease. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, 154:12–24.

Falaki, A., Huang, X., Lewis, M. M., and Latash, M. L. (2017). Dopaminergic modulation of multi-muscle synergies in postural tasks performed by patients with parkinson's disease. *Journal of Electromyography and Kinesiology: Official Journal of the International Society of Electrophysiological Kinesiology*, 33:20–26.

Falaki, A., Jo, H. J., Lewis, M. M., O'Connell, B., De Jesus, S., McInerney, J., Huang, X., and Latash, M. L. (2018). Systemic effects of deep brain stimulation on synergic control in parkinson's disease. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*.

Fang, H.-S., Li, J., Tang, H., Xu, C., Zhu, H., Xiu, Y., Li, Y.-L., and Lu, C. (2022). Alphapose: Whole-body regional multi-person pose estimation and tracking in real-time. *arXiv preprint arXiv:2211.03375*. Accepted to IEEE TPAMI.

Feigin, V. L., Stark, B. A., Johnson, C. O., Roth, G. A., Bisignano, C., Abady, G. G., Abbasifard, M., Abbasi-Kangevari, M., Abd-Allah, F., Abdelalim, A., and et al. (2021). Global, regional, and national burden of stroke and its risk factors, 1990–2019: a systematic analysis for the global burden of disease study 2019. *The Lancet Neurology*, 20(10):795–820.

Fischer, S., Hallschmid, M., Elsner, A., and Born, J. (2002). Sleep forms memory for finger skills. *Proceedings of the National Academy of Sciences*, 99:11987–11991.

Fischer, S., Nitschke, M. F., Melchert, U. H., Erdmann, C., and Born, J. (2005). Motor memory consolidation in sleep shapes more effective neuronal representations. *Journal of Neuroscience*, 25(49):11248–11255.

Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47(6):381–391.

Fitts, P. M. (1964). Perceptual-motor skill learning. *Categories of human learning*, 47:381–391.

Fitts, P. M. and Posner, M. I. (1967). *Human Performance*. Brooks/Cole.

Flanagan, J. R. and Wing, A. M. (1997). The role of internal models in motion planning and control: Evidence from grip force adjustments during movements of hand-held loads. *Journal of Neuroscience*, 17(4):1519–1528.

Franklin, D. W. and Wolpert, D. M. (2008). Specificity of reflex adaptation for task-relevant variability. *Journal of Neuroscience*, 28(52):14165–14175.

Fricke, C., Gentner, R., Alizadeh, J., and Classen, J. (2020). Linking individual movements to a skilled repertoire: fast modulation of motor synergies by repetition of stereotyped movements. *Cerebral Cortex*, 30(3):1185–1198.

Gelman, A. (2003). A bayesian formulation of exploratory data analysis and goodness-of-fit testing. *International Statistical Review*, 71(2):369–382.

Gentile, A. M. (2000). *Skill acquisition: Action, movement, and the neuromotor processes*. FA Davis Company.

Gentner, R., Gorges, S., Weise, D., aufm Kampe, K., Buttmann, M., and Classen, J. (2010). Encoding of motor skill in the corticomuscular system of musicians. *Current Biology*, 20(20):1869–1874.

Gobet, F. and Simon, H. A. (1996). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology*, 31(1):1–40.

Gonzalez Castro, L. N., Hadjiosif, A. M., Hemphill, M. A., and Smith, M. A. (2014). Environmental consistency determines the rate of motor adaptation. *Current Biology*, 24(10):1050–1061.

Gorniak, S. L., Zatsiorsky, V. M., and Latash, M. L. (2007). Hierarchies of synergies: an example of the two-hand, multi-finger tasks. *Experimental Brain Research*, 179(2):167–180.

Gorniak, S. L., Zatsiorsky, V. M., and Latash, M. L. (2009). Hierarchical control of static prehension: ii. multi-digit synergies. *Experimental Brain Research*, 194(1):1–15.

Graybiel, A. M. (1998). The basal ganglia and chunking of action repertoires. *Neurobiology of Learning and Memory*, 70(1-2):119–136.

Graybiel, A. M. (2008). Habits, rituals, and the evaluative brain. *Annual Review of Neuroscience*, 31:359–387.

Grefkes, C. and Fink, G. R. (2011). Reorganization of cerebral networks after stroke: new insights from neuroimaging with connectivity approaches. *Brain*, 134(5):1264–1276.

Güler, R. A., Neverova, N., and Kokkinos, I. (2018). Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7297–7306.

Haar, S. and Faisal, A. A. (2020). Brain activity reveals multiple motor-learning mechanisms in a real-world task. *Frontiers in Human Neuroscience*, 14:354.

Haith, A. M. and Krakauer, J. W. (2018). The multiple effects of practice: skill, habit and reduced cognitive load. *Current Opinion in Behavioral Sciences*, 20:196–201. Published in final edited form as: Curr Opin Behav Sci. 2018 Mar 1;20:196–201.

Hardwick, R. M., Caspers, S., Eickhoff, S. B., and Swinnen, S. P. (2018). Neural correlates of action: Comparing meta-analyses of imagery, observation, and execution. *Neuroscience & Biobehavioral Reviews*, 94:31–44.

Hardwick, R. M., Rottschy, C., Miall, R. C., and Eickhoff, S. B. (2013). A quantitative meta-analysis and review of motor learning in the human brain. *NeuroImage*, 67:283–297.

Harris, C. M. and Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394:780–784.

Hastie, T., Tibshirani, R., and Friedman, J. (2011). *The Elements of Statistical Learning*. Springer.

Hausmann, S. B., Vargas, A. M., Mathis, A., and Mathis, M. W. (2021). Measuring and modeling the motor system with machine learning. *arXiv preprint arXiv:2103.11775*.

Herzog, J., Schweighofer, N., and Valero-Cuevas, F. J. (2024). Changes in muscle synergy structure and activation patterns underlie force field adaptation, retention, and generalization. *bioRxiv*.

Hikosaka, O., Nakamura, K., Sakai, K., and Nakahara, H. (2002). Central mechanisms of motor skill learning. *Current Opinion in Neurobiology*, 12(2):217–222.

Hoehn, M. M. and Yahr, M. D. (1967). Parkinsonism: Onset, progression and mortality. *Neurology*, 17:427–442.

Hong, Y.-N. G., Ballekere, A. N., Fregly, B. J., and Roh, J. (2021). Are muscle synergies useful for stroke rehabilitation? *Current Opinion in Biomedical Engineering*, 19:100315.

Hotermans, C., Peigneux, P., Maertens de Noordhout, A., Moonen, G., and Maquet, P. (2006). Early boost and slow consolidation in motor skill learning. *Learning & Memory*, 13(5):580–583.

Houk, J. C., Buckingham, J. T., and Barto, A. G. (1996). Models of the cerebellum and motor learning. *Behavioral and Brain Sciences*, 19(3):368–383.

Hsu, A. I. and Yttri, E. A. (2021). B-soid: An open-source unsupervised algorithm for identification and fast prediction of behaviors. *Nature Communications*, 12(1):1–13.

Hughlings Jackson, J. (1889). On the comparative study of diseases of the nervous system. *Lancet*, pages 617–623.

Hunt, R. H. and Aslin, R. N. (2001). Statistical learning in a serial reaction time task: Access to separable statistical cues by individual learners. *Journal of Experimental Psychology: General*, 130(4):658–680.

Insafutdinov, E., Pishchulin, L., Andres, B., Andriluka, M., and Schiele, B. (2017). Arttrack: Articulated multi-person tracking in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6457–6465.

Jain, A. K. (2010). Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8):651–666.

Jeannerod, M. (1988). *The neural and behavioural organization of goal-directed movements*. Oxford University Press, New York.

Jo, H. J., Maenza, C., Good, D. C., Huang, X., Park, J., Sainburg, R. L., and Latash, M. L. (2016a). Effects of unilateral stroke on multi-finger synergies and their feed-forward adjustments. *Neuroscience*.

Jo, H. J., Mattos, D., Lucassen, E. B., Huang, X., and Latash, M. L. (2016b). Changes in multidigit synergies and their feed-forward adjustments in multiple sclerosis. *Journal of Motor Behavior*, pages 1–11.

Johansson, H., Folkerts, A.-K., Hammarström, I., Kalbe, E., and Leavy, B. (2023). Effects of motor-cognitive training on dual-task performance in people with parkinson's disease: a systematic review and meta-analysis. *Journal of Neurology*, 270(6):2890–2907.

Johnson, N. F. (1970). The role of chunking and organization in the process of recall. In Bower, G. H., editor, *Psychology of Learning and Motivation*, volume 4, pages 171–247. Academic Press.

Kalaska, J. F. and Crammond, D. J. (1992). Cerebral cortical mechanisms of reaching movements. *Science*, 255(5051):1517–1523.

Kantak, S. S., Sullivan, K. J., Fisher, B. E., Knowlton, B. J., and Winstein, C. J. (2010). Neural substrates of motor memory consolidation depend on practice structure. *Nature Neuroscience*, 13(8):923–925.

Karni, A., Meyer, G., Jezzard, P., Adams, M. M., Turner, R., and Ungerleider, L. G. (1995). Functional mri evidence for adult motor cortex plasticity during motor skill learning. *Nature*, 377:155–158.

Kaufmann, P., Koller, W., Wallnöfer, E., Stöllinger, T., Kogler, G., Pellegrini, G., Hasson, C. J., Riener, R., Wieser, S., Ilg, W., and Kuster, A. D. (2024). Increased trial-to-trial similarity and reduced temporal overlap of muscle synergy activation coefficients manifest during learning and with increasing movement proficiency. *Scientific Reports*, 14:17638.

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9:718–727.

Kendall, A., Grimes, M., and Cipolla, R. (2015). Posenet: A convolutional network

for real-time 6-dof camera relocalization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2938–2946.

Kennerley, S. W., Sakai, K., and Rushworth, M. F. S. (2004). Organization of action sequences and the role of the pre-sma. *Journal of Neurophysiology*, 91(2):978–993.

Kitago, T. and Krakauer, J. W. (2016). Ten guiding principles for movement training in neurorehabilitation. *Archives of Physiotherapy*, 6:1–11.

Kleim, J. A. (2011). Neural plasticity and neurorehabilitation: Teaching the new brain old tricks. *Journal of Communication Disorders*, 44(5):521–528.

Komar, J., Seifert, L., Vergne, N., Simon, I. N., Lardy, B., Chow, J. W., Mourot, L., and Slawinski, J. (2023). Narrowing the coordination solution space during motor learning standardizes individual patterns of search strategy but diversifies learning rates. *Scientific Reports*, 13:2009.

Kording, K. P. and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247.

Kornysheva, K., Bush, D., Meyer, S. S., Sadnicka, A., Barnes, G., and Burgess, N. (2019). Neural competitive queuing of ordinal structure underlies skilled sequential action. *Neuron*, 101(6):1166–1180.e3.

Kornysheva, K. and Diedrichsen, J. (2014). Human premotor areas parse sequences into their spatial and temporal features. *eLife*, 3:e03043.

Koziol, L. F., Budding, D., Andreasen, N., D'Arrigo, S., Bulgheroni, S., Imamizu, H., Ito, M., Manto, M., Marvel, C., Parker, K., Pezzulo, G., Ramnani, N., Riva, D., Schmahmann, J., Vandervert, L., and Yamazaki, T. (2014). Consensus paper: The cerebellum's role in movement and cognition. *Cerebellum*, 13(1):151–177.

Krakauer, J. W. (2009). Motor learning and consolidation: The case of visuomotor rotation. *Advances in Experimental Medicine and Biology*, 629:405–421.

Krakauer, J. W. and Mazzoni, P. (2011). Human sensorimotor learning: adaptation, skill, and beyond. *Current Opinion in Neurobiology*, 21(4):636–644. Epub 2011 Jul 20.

Krakauer, J. W. and Mazzoni, P. (2019). *Motor Learning and Control: From Theory to Practice*. Academic Press.

Krakauer, J. W., Pine, Z. M., Ghilardi, M. F., and Ghez, C. (2000). Learning of visuomotor transformations for vectorial planning of reaching trajectories. *Journal of Neuroscience*, 20(23):8916–8924.

Krakauer, J. W. and Shadmehr, R. (2006). Consolidation of motor memory. *Trends in Neurosciences*, 29:58–64.

Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86.

Kutsuzawa, T. and Hayashibe, M. (2022). Motor synergy generalization framework for new targets in multi-planar and multi-directional reaching tasks. *Royal Society Open Science*, 9(12):211721.

Laird, J. E., Rosenbloom, P. S., and Newell, A. (1984). Towards chunking as a general learning mechanism. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-84)*, pages 188–192.

Lang, C. E., Bland, M. D., Bailey, R. R., Schaefer, S. Y., and Birkenmeier, R. L. (2013). Assessment of upper extremity impairment, function, and activity after stroke: foundations for clinical decision making. *Journal of Hand Therapy*, 26(2):104–114. quiz 115.

Lashley, K. S. (1951). The problem of serial order in behavior. In Jeffress, L. A., editor, *Cerebral Mechanisms in Behavior*, pages 112–136. Wiley.

Latash, M. L. (2008a). *Fundamentals of Motor Control*. Academic Press.

Latash, M. L. (2008b). *Synergies: Neural Organization and Functional Integration*. Oxford University Press.

Latash, M. L. (2010). Stages in learning motor synergies: a view based on the equilibrium-point hypothesis. *Human Movement Science*, 29(5):642–654. Epub 2010 Jan 8.

Latash, M. L. (2020). Synergies: A theoretical approach to understanding movement control. *Motor Control*, 24(3):392–414.

Latash, M. L. (2021). One more time about motor (and non-motor) synergies. *Experimental Brain Research*, 239(10):2951–2967.

Latash, M. L. and Huang, X. (2015). Neural control of movement stability: Synergies, subcortical loops, and impairments. *Neuroscientist*, 21:69–76.

Latash, M. L., Scholz, J. P., and Schöner, G. (2007). Toward a new theory of motor synergies. *Motor Control*, 11:276–308.

Lee, D. D. and Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791.

Levin, M. F. and Demers, M. (2021). Motor learning in neurological rehabilitation. *Disability and Rehabilitation*, 43(24):3445–3453.

Levin, M. F., Michaelsen, S. M., Cirstea, C. M., and Roby-Brami, A. (2002). Use of the trunk for reaching targets placed within and beyond the reach in adult hemiparesis. *Experimental Brain Research*, 143(2):171–180.

Lin, J. (1991). Divergence measures based on the shannon entropy. *IEEE Transactions on Information Theory*, 37(1):145–151.

Linderman, S., Johnson, M., Miller, A., Adams, R., Blei, D., and Paninski, L. (2017). Bayesian learning and inference in recurrent switching linear dynamical systems. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, PMLR*, pages 914–922.

Lindsay, A. J. and Seamans, J. K. (2024). HUB-DT: A tool for unsupervised behavioural discovery and analysis. *bioRxiv*, page 2024.01.08.574696.

Lugaresi, C., Tang, J., Nash, H., McClanahan, C., Uboweja, E., Hays, M., Zhang, F., Chang, C.-L., Yong, M. G., Lee, J., Chang, W.-T., Hua, W., Georg, M., and Grundmann, M. (2019). Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*.

Luxem, K., Mocellin, P., Fuhrmann, F., Kürsch, J., and Miller, S. R. (2020). Identifying behavioral structure from deep variational embeddings of animal motion. *bioRxiv*.

Madarshahian, S. and Latash, M. L. (2021). Synergies at the level of motor units in single-finger and multi-finger tasks. *Experimental Brain Research*, 239(4):961–973.

Maier, M., Ballester, B. R., and Verschure, P. F. M. J. (2019). Principles of neurorehabilitation after stroke based on motor learning and brain plasticity mechanisms. *Frontiers in Systems Neuroscience*, 13:74.

Manto, M., Bower, J. M., Conforto, A. B., Delgado-García, J. M., da Guarda, S. N. F., Gerwig, M., Habas, C., Hagura, N., Ivry, R. B., and Mariën, P. (2012). The cerebellum: a hub for balancing motor and cognitive functions. *Nature Reviews Neuroscience*, 13(6):360–368.

Marconi, R., Landi, A., and Valzania, F. (2008). Subthalamic nucleus stimulation in parkinson's disease. *Neurological Sciences*, 29(5):S389–S391.

Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., and Bethge, M. (2018). Deeplabcut: Markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 21(9):1281–1289.

McGowan, C. P., Neptune, R. R., Clark, D. J., and Kautz, S. A. (2010). Modular control of human walking: Adaptations to altered mechanical demands. *Journal of Biomechanics*, 43(3):412–419.

McInnes, L., Healy, J., and Astels, S. (2017). hdbscan: Hierarchical density based clustering. *Journal of Open Source Software*, 2(11):205.

McInnes, L., Healy, J., and Melville, J. (2018). Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv*, 1802.03426.

McPherson, L. M. and Dewald, J. P. A. (2022). Abnormal synergies and associated reactions post-hemiparetic stroke reflect muscle activation patterns of brainstem motor pathways. *Frontiers in Neurology*, 13.

Meyendorff, J. (1974). *A Study of Gregory Palamas*. St. Vladimir's Seminary Press.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2):81–97.

Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research*, 42:223–227.

Morehead, J. R. and Orban de Xivry, J.-J. (2021). A synthesis of the many errors and learning processes of visuomotor adaptation. *bioRxiv*, page 2021.03.14.435278.

Mussa-Ivaldi, F. A. (1999). Modular features of motor control and learning. *Current Opinion in Neurobiology*, 9(6):713–717.

Nader, K., Schafe, G. E., and LeDoux, J. E. (2000). The labile nature of consolidation theory. *Nature Reviews Neuroscience*, 1:216–222.

Newell, K. M. (1991). *Motor learning: Concepts and applications*. Macmillan.

Nichols-Larsen, D. S., Clark, P. C., Zeringue, A., Greenspan, A., and Blanton, S. (2005). Factors influencing stroke survivors' quality of life during subacute recovery. *Stroke*, 36(7):1480–1484.

Nishida, M. and Walker, M. P. (2007). Daytime naps, motor memory consolidation and regionally specific sleep spindles. *PLoS ONE*, 2(4):e341.

Nissen, M. J. and Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19:1–32.

Overduin, S. A., d'Avella, A., Carmena, J. M., and Bizzi, E. (2012). Microstimulation activates a handful of muscle synergies. *Neuron*, 76(6):1071–1077.

Overduin, S. A., d'Avella, A., Roh, J., and Bizzi, E. (2008). Modulation of muscle synergy recruitment in primate grasping. *Journal of Neuroscience*, 28(4):880–892.

Palamas, G. (1983). *The Triads*. Paulist Press.

Palamas, G. (1988). *The One Hundred and Fifty Chapters*. Paulist Press.

Park, J.-K., Lewis, M. M., Huang, X., and Latash, M. L. (2014). Dopaminergic modulation of motor coordination in parkinson's disease. *Parkinsonism & Related Disorders*, 20:64–68.

Penhune, V. B. and Steele, C. J. (2012). Parallel contributions of cerebellar, striatal, and m1 mechanisms to motor sequence learning. *Behavioural Brain Research*, 226(2):579–591.

Perruchet, P. and Pacton, S. (2006). Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in Cognitive Sciences*, 10(5):233–238.

Poggio, T. and Bizzi, E. (2004). Generalization in vision and motor control. *Nature*, 431(7010):768–774.

Poldrack, R. A., Sabb, F. W., Foerde, K., Tom, S. M., Asarnow, R. F., Bookheimer, S. Y., and Knowlton, B. J. (2005). The neural correlates of motor skill automaticity. *Journal of Neuroscience*, 25(22):5356–5364.

Ramat, S., Leigh, R. J., Zee, D. S., and Optican, L. M. (2007). What clinical disorders tell us about the neural control of saccadic eye movements. *Brain*, 130(1):10–35.

Ramkumar, P., Acuna, D., Berniker, M., Grafton, S., Turner, R. S., and Kording, K. (2016). Chunking as the result of an efficiency computation trade-off. *Nature Communications*, 7:12176.

Reisman, D. S. and Scholz, J. P. (2003). Aspects of joint coordination are preserved during pointing in persons with post-stroke hemiparesis. *Brain*, 126(Pt 11):2510–2527.

Rioult-Pedotti, M. S., Friedman, D., and Donoghue, J. P. (2000). Learning-induced ltp in neocortex. *Science*, 290(5491):533–536.

Robertson, E. M. (2005). Off-line learning and the primary motor cortex. *Journal of Neuroscience*, 25(27):6372–6378.

Robertson, E. M. (2019). Skill memory: Mind the ever-decreasing gap for offline processing. *Current Biology*, 29(8):R287–R289.

Robinson, D. A. (2022). Neurophysiology, pathology and models of rapid eye movements. In *Progress in Brain Research*, volume 267, chapter 14, pages 287–317. Elsevier.

Rocchi, L., Carlson-Kuhta, P., Chiari, L., Burchiel, K. J., Hogarth, P., and Horak, F. B. (2012). Effects of deep brain stimulation in the subthalamic nucleus or globus pallidus internus on step initiation in parkinson's disease: Laboratory investigation. *Journal of Neurosurgery*, 117(6):1141–1149.

Salbach, N. M., Mayo, N. E., Robichaud-Ekstrand, S., Hanley, J. A., Richards, C. L., and Wood-Dauphinee, S. (2006). Balance self-efficacy and its relevance to physical function and perceived health status after stroke. *Archives of Physical Medicine and Rehabilitation*, 87(3):364–370.

Sanger, T. D. (2000). Human arm movements described by a low-dimensional superposition of principal components. *Journal of Neuroscience*, 20(3):1066–1072.

Santello, M. (2002). Kinematic synergies for the control of hand shape. *Archives Italiennes de Biologie*, 140(3):221–228.

Santello, M., Baud-Bovy, G., and Jörntell, H. (2016). Hand synergies: integration of robotics and neuroscience for understanding the control of biological and artificial hands. *Physics of Life Reviews*, 17:1–23.

Sarwary, A. M. E., Stegeman, D. F., Selen, L. P. J., and Medendorp, W. P. (2015). Generalization and transfer of contextual cues in motor learning. *Journal of Neurophysiology*, 114(1):479–488.

Schapiro, A. C., McDevitt, E. A., Rogers, T. T., Mednick, S. C., Norman, K. A., and Brawn, T. P. (2018). Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nature Communications*, 8:1131.

Schmidt, R. A. (1975). *A schema theory of discrete motor skill learning*. Psychological review.

Schmidt, R. A. and Lee, T. D. (2005). *Motor Control and Learning: A Behavioral Emphasis*. Human Kinetics, 5th edition.

Schoenfeld, M. J., Thom, J., Williams, J., Stagg, C. J., and Zich, C. (2024). Relationship between skill training and skill transfer through the example of bimanual motor learning. *European Journal of Neuroscience*, 59(1):54–68.

Scholz, J. P. and Schöner, G. (1999). Uncontrolled manifold concept: identifying control variables for a functional task. *Experimental Brain Research*, 126:289–306.

Scott, S. H. (2004). Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5(7):532–546.

Seidler, R. D., Bernard, J. A., Burutolu, T. B., Fling, B. W., Gordon, M. T., Gwin, J. T., Kwak, Y., and Lipps, D. B. (2010). Motor control and aging: Links to age-related brain structural, functional, and biochemical effects. *Neuroscience and Biobehavioral Reviews*, 34:721–733.

Servan-Schreiber, D. and Anderson, J. R. (1990). Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16(4):592–608.

Shadmehr, R. and Krakauer, J. W. (2010). A computational neuroanatomy for motor control. *Experimental Brain Research*, 207:319–336.

Shadmehr, R. and Mussa-Ivaldi, S. (2012). *Biological Learning and Control: How the Brain Builds Representations, Predicts Events, and Makes Decisions*. MIT Press, Cambridge, MA.

Shadmehr, R. and Wise, S. P. (2005). *The Computational Neurobiology of Reaching and Pointing*. MIT Press.

Singhal, A. (2001). Modern information retrieval: A brief overview. *IEEE Data Engineering Bulletin*, 24(4):35–43.

Smith, C. (2001). Sleep states and memory processes in humans: Procedural versus declarative memory systems. *Sleep Medicine Reviews*, 5:491–506.

Song, S. and Cohen, L. G. (2014). Practice and sleep form different aspects of skill. *Nature Communications*, 5:3407.

Stagg, C. J. and Johansen-Berg, H. (2013). Studying the effects of transcranial direct-current stimulation in stroke recovery using magnetic resonance imaging. *Frontiers in Human Neuroscience*, 7:857.

Stickgold, R. and Walker, M. P. (2007). Sleep-dependent memory consolidation and reconsolidation. *Sleep Medicine*, 8(4):331–343.

Stinear, C. M., Petoe, M. A., and Byblow, W. D. (2015). Primary motor cortex excitability during recovery after stroke: implications for neuromodulation. *Brain Stimulation*, 8(6):1183–1190.

Tamaki, M., Huang, T. R., Yotsumoto, Y., and Hämäläinen, M. S. (2013). Enhanced spontaneous oscillations in the supplementary motor area are associated with

sleep-dependent offline learning of finger-tapping motor-sequence task. *Journal of Neuroscience*, 33(34):13894–13902.

Tanji, J. and Shima, K. (1994). Role for supplementary motor area cells in planning several movements ahead. *Nature*, 371(6496):413–416.

Taylor, J. A. and Ivry, R. B. (2013). Implicit and explicit processes in motor learning. In Prinz, W., Beisert, M., and Herwig, A., editors, *Action Science: Foundations of an Emerging Discipline*, pages 63–88. MIT Press, Cambridge, MA.

Taylor, J. A., Krakauer, J. W., and Ivry, R. B. (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience*, 34(8):3023–3032.

Thach, W. T., Goodkin, H. P., and Keating, J. G. (1992). The cerebellum and the adaptive coordination of movement. *Annual Review of Neuroscience*, 15:403–442.

Ting, L. H. and McKay, J. L. (2008). Neuromechanics of muscle synergies for posture and movement. *Current Opinion in Neurobiology*, 17(6):622–628.

Todorov, E. and Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5:1226–1235.

Tomita, Y., Turpin, N. A., Piscitelli, D., Feldman, A. G., and Levin, M. F. (2020). Stability of reaching during standing in stroke. *Journal of Neurophysiology*, 123(5):1756–1765.

Tresch, M. C., Cheung, V. C., and d'Avella, A. (2006). Matrix factorization algorithms for the identification of muscle synergies: evaluation on simulated and experimental data sets. *Journal of Neurophysiology*, 95(4):2199–2212.

Tresch, M. C. and Jarc, A. (2009). The case for and against muscle synergies. *Current Opinion in Neurobiology*, 19(6):601–607.

Tresch, M. C., Saltiel, P., and Bizzi, E. (1999). The construction of movement by the spinal cord. *Nature Neuroscience*, 2(2):162–167.

Turvey, M. T. (1990). Coordination. *American Psychologist*, 45(8):938–953.

Turvey, M. T. (2007). Action and perception at the level of synergies. *Human Movement Science*, 26(4):657–697.

Törnbom, K., Lundälv, J., and Sunnerhagen, K. S. (2019). Long-term participation 7–8 years after stroke: Experiences of people in working-age. *PLOS ONE*, 14(3):e0213447.

Valero-Cuevas, F. J. (2019). Dimensionality of joint torques during continuous, cyclic movements: Using linear and nonlinear dimensionality reduction methods to study coordination. *PLOS ONE*, 14:e0214742.

van der Maaten, L. J. P. (2014). Accelerating t-sne using tree-based algorithms. *Journal of Machine Learning Research*, 15:3221–3245.

van der Maaten, L. J. P. and Hinton, G. E. (2008). Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605.

Verwey, W. B. (1996). Buffer loading and chunking in sequential keypressing. *Journal of Experimental Psychology: Human Perception and Performance*, 22:544–562.

Verwey, W. B. (2001). Concatenating familiar movement sequences: The versatile cognitive processor. *Acta Psychologica*, 106(1-2):69–95.

Verwey, W. B. and Eikelboom, T. (2003). Evidence for lasting sequence segmentation in the discrete sequence-production task. *Journal of Motor Behavior*, 35(2):171–181.

Verwey, W. B., Shea, C. H., and Wright, D. L. (2015). A cognitive framework for explaining serial processing and sequence execution. *Psychonomic Bulletin & Review*, 22(1):54–77.

Villani, C. (2008). *Optimal transport: old and new*, volume 338. Springer Science & Business Media.

Wagner, F. B., Kim, J., and Schnitzer, M. J. (2020). Skilled reaching tasks for head-fixed mice using a robotic manipulandum. *Nature Protocols*, 15:2699–2714.

Walker, M. P., Brakefield, T., Morgan, A., Hobson, J. A., and Stickgold, R. (2002). Practice with sleep makes perfect: Sleep-dependent motor skill learning. *Neuron*, 35:205–211.

Walker, M. P. and Stickgold, R. (2003). Sleep-dependent learning and memory consolidation. *Neuron*, 44:121–133.

Walker, M. P. and Stickgold, R. (2005). It's practice, with sleep, that makes perfect: Implications of sleep-dependent learning and plasticity for skill performance. *Clinical Sports Medicine*, 24:301–317.

Wang, J. and Sainburg, R. L. (2005). Adaptation to visuomotor rotations remaps movement vectors, not final positions. *Journal of Neuroscience*, 25(15):4024–4030.

Wang, J. and Sainburg, R. L. (2007). The dominant and nondominant arms are specialized for stabilizing different features of task performance. *Experimental Brain Research*, 178(4):565–570.

Ward, N. S., Brander, F., and Kelly, K. (2019). Intensive upper limb neurorehabilitation in chronic stroke: outcomes from the queen square programme. *Journal of Neurology, Neurosurgery & Psychiatry*, 90(5):498–506.

Weaver, I., Roach, J., Dunn, T., and Tadross, M. (2024). System and method for three-dimensional kinematic profiling of animal models.

Welsh, J. P. and Llinas, R. (1997). Some organizing principles for the control of movement based on olivocerebellar physiology. *Progress in Brain Research*, 114:449–461.

Wiltschko, A. B., Johnson, M. J., Iurilli, G., Peterson, R. E., Katon, J. M., Pashkovski, S. L., Abraira, V. E., Adams, R. P., and Datta, S. R. (2015). Mapping sub-second structure in mouse behavior. *Neuron*, 88(6):1121–1135.

Winter, D. A. (1990). *Biomechanics and Motor Control of Human Movement*. John Wiley & Sons, New York, 2nd edition.

Wolf, S. L., Winstein, C. J., Miller, J. P., Taub, E., Uswatte, G., Morris, D., Giuliani, C., Light, K. E., and Nichols-Larsen, D. (2006). Effect of constraint-induced movement therapy on upper extremity function 3 to 9 months after stroke: The excite randomized clinical trial. *JAMA*, 296(17):2095–2104.

Wolpert, D. M., Diedrichsen, J., and Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature Reviews Neuroscience*, 12(12):739–751.

Wolpert, D. M. and Flanagan, J. R. (2001). Motor prediction. *Current Biology*, 11(18):R729–R732.

Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269:1880–1882.

Wong, K. C. (2019). A tutorial on variational bayesian inference. *arXiv preprint arXiv:1912.10776*.

Wulf, G. (2008). Attentional focus effects in balance acrobats. *Research Quarterly for Exercise and Sport*, 79(3):319–325.

Wulf, G. and Lewthwaite, R. (2010). Effortless motor learning? an external focus of attention enhances movement effectiveness and efficiency. *Current Directions in Psychological Science*, 19(2):59–64.

Wymbs, N. F., Bassett, D. S., Mucha, P. J., Porter, M. A., and Grafton, S. T. (2012). Differential recruitment of the sensorimotor putamen and frontoparietal cortex during motor chunking in humans. *Neuron*, 74(5):936–946.

Xu, R. and Wunsch, D. (2005). Survey of clustering algorithms. *IEEE Transactions on Neural Networks*, 16(3):645–678.

Xu, R. and Wunsch, D. C. (2003). *Clustering*. Wiley-IEEE Press.

Yokoyama, H., Kaneko, N., Ogawa, T., Kawashima, N., Watanabe, K., and Nakazawa, K. (2019). Cortical correlates of locomotor muscle synergy activation in humans: An electroencephalographic decoding study. *iScience*, 15:623–639.

Zhao, K., He, C., Xiang, W., Zhou, Y., Zhang, Z., Li, J., and Scano, A. (2023). Evidence of synergy coordination patterns of upper-limb motor control in stroke patients with mild and moderate impairment. *Frontiers in Physiology*, 14:1214995.

Zimnik, A. J. and Churchland, M. M. (2021). Independent generation of sequence elements by motor cortex. *Nature Neuroscience*, 24(3):412–424.