

Machine learning models and their development process as learning affordances for humans

Carmel Kent¹, Muhammad Ali Chaudhry², Mutlu Cukurova², Ibrahim Bashir¹,
Hannah Pickard³, Chris Jenkins³, Benedict du Boulay⁴, Anissa Moeini¹ and Rosemary
Luckin²

¹ EDUCATE Ventures

² UCL Knowledge Lab, University College London

³ ZISHI, OSTC Group

⁴ University of Sussex

Abstract. This paper explores the relationship between unsupervised machine learning models, and the mental models of those who develop or use them. In particular, we consider unsupervised models, as well as the 'organisational co-learning process' that creates them, as learning affordances. The co-learning process involves inputs originating both from the human participants' shared semantics, as well as from the data. By combining these, the process as well as the resulting computational models afford a newly shaped mental model, which is potentially more resistant to the biases of human mental models. We illustrate this organisational co-learning process with a case study involving unsupervised modelling via commonly used methods such as dimension reduction and clustering. Our case study describes how a trading and training company engaged in the co-learning process, and how its mental models of trading behavior were shaped (and afforded) by the resulting unsupervised machine learning model. The paper argues that this kind of co-learning process can play a significant role in human learning, by shaping and safeguarding participants' mental models, precisely because the models are unsupervised, and thus potentially lead to learning from unexpected or inexplicit patterns.

Keywords: Learners' mental models, unsupervised machine learning, co-learning process.

1 Introduction

It is well established in the learning literature that presenting learners with a simplified model of whatever is to be understood is a helpful step in learning [1]. One example of these simplified models is the use of "notional machines" in teaching about programs, computers and programming [2]. The programming teacher offers analogies such as "a variable is like a box" or possibly draws a simplified diagram of how a loop in a program works. Of course, learners' consequent mental models will not normally exactly match the simplified model that was presented, but they are normally influenced by it, and such simplified models also help towards a better understanding of the more complex truth from which the simplified model has been derived. The use of simplified models can assist organisations as well as individuals in learning new material, as well

as about themselves. In the case of managers engaging in self-understanding of their organisation [3], machine learning (ML) techniques applied to organisational data can act as a mirror back to them and thus to the organisation [4]. In this respect, unsupervised learning methods play a particular role as they are more likely to reveal factors that the managers were not explicitly aware of. This paper extends on a preliminary work [5], and its main contribution is not in the modelling itself, but rather in the theoretical postulation that ML models, generated by unsupervised methods, can be used as learning affordances to support the development of the mental models of the managers in an organisation. For clarity in what follows, we are not concerned here with how the managers' changed mental models diffused through the organisation, but with the fact that something changed in the mental models of those managers.

The ZISHI/OSTC company trades in “futures”¹ and is a training company largely for university graduates, who join the company to learn the art of trading. The senior training managers who were leading this work in ZISHI, and were also interviewed in this study, are interchangeably referred to throughout this paper as either managers, trainers or experts. In order for them to support their learners (during and after two months of formal training), and to design mentoring/training tools for them, they needed first to understand what did trading actually look like in their own context, after years of nurturing tacit mental models. The trainers certainly had a strong sense that different traders traded in different ways and had developed a partial typology of trading behaviors: for example, some traders preferred to work in volatile markets, others in more stable markets. Based on this implicit mental model, trainers might suggest different markets to individual traders based on this preference. However, the typology had not been reified within the company and had remained largely tacit.

In an attempt to define ‘learning’, [6] (adapted from [7]) states that learning is “*a process that leads to change, which occurs as a result of experience and increases the potential for improved performance and future learning*” (p.3). Therefore, in order to help the company better understand its traders' behavior and maintain a culture of learning and change [8], we used unsupervised ML methods to arrive at four multidimensional profiles of trading behavior. In parallel, we asked the company to generate its own, till then largely tacit [9], trading behavior profiles into written descriptions. We were then able to compare these data-driven profiles with the company's self-generated profiling. After comparison and validation, the data-driven profiles (being validated as a refinement of the original mental profiling model) were used as the basis of a predictive decision support tools for hiring and mentoring, both tools are out of the scope of this paper.

We term the process of crafting the four behavioral profiles using ML methods a ‘co-learning process’, since this is a process combined of human ‘supervision’ in some of its stages, where the experts' semantics guides the analysis, and is unsupervised in other stages, where the modeling is being guided purely by the data patterns. In that sense, it embodies some of the notions of connectivism [10] or extended minds [11], where learning is not considered as residing within a single learner, but rather stems

¹ Futures are derivative financial contracts that obligate the parties to transact an asset at a predetermined future date and price.

from the continuous interactions between learners, their organisation, and artefacts such as data.

This paper is organised into five sections. The next section focuses on the distinction between unsupervised ML models and mental models. Section 3 describes the ways in which unsupervised ML methods were used as potential learning affordances. Section 4 compares the resulting ML model with the company's original mental model. Section 5 concludes and summarises our findings.

2 Unsupervised ML models and mental models

Raw data are not independent, contextless, self-sufficient repositories of meaning [12]. Contextualised modeling of data, using statistical methods and, particularly, ML, creates possibilities for assigning existing semantics to the models, as well as for creating new semantics, which in turn, can be used as **“learning affordances”**. The concept of affordance describes the complementary relationship between an environment and what it offers or provides to the actors within it [13]. The process of data modeling, which we refer to here as a “learning affordance”, can start from a phase of feature engineering, in which the existing semantics can be attached to the raw data to shape it in a contextualised way. Later the process can generate (or rather bring to the surface) new hidden or implicit meanings, using methods such as unsupervised ML. In the unsupervised learning phase, hidden statistical relationships, or other statistical constructs (such as distributions) will emerge, to be interpreted via the stakeholders' original mental models (which are based on expectations, projections, cognitive biases, and emotions), as well as generating new inferences and new assumptions. This can be an iterative process, in which the unsupervised model will be reviewed, new hypotheses raised, and the model tweaked and refined to serve an augmented purpose. Its purpose, as we will suggest in this paper, can be to support individual or organisational learning, by externalising and simplifying existing mental models of the world they are learning about.

This paper argues that the model generated by unsupervised methods provides a learning affordance, not just because it simplifies, corrects and highlights different aspects of an existing mental model, but also because it can enable the creation of a new semantics and a new language to revise that mental model. The ML model is generated by a process, which by itself can be considered as a learning affordance in the sense that it offers a useful dialogical entity between knowledge existing in human minds, and the patterns arising from the data, and by means of that – might cancel out mutual biases, and open new opportunities for learning.

In many senses, supervised ML and reinforcement algorithms inherently include in them the aspiration to mimic some specific human behavior and to optimise on the basis of human observations. Unsupervised learning, on the other hand, can reveal factors and behaviors that human guidance might have been preventing us from seeing. In an analogy to human learning, a child might learn purely from observing (even if the scene is to some extent orchestrated by an adult human), and not always by following an adult deliberately pointing out (i.e., supervising), rewarding or punishing (i.e., reinforcing) to guide her. Unsupervised ML algorithms, such as clustering, dimension reduction or association techniques, are designed without a top-down supervision component, and

in many respects, the only human intervention will be expressed prior to the algorithm itself, in the form of feature engineering. Thus, unsupervised algorithms are more about identification than recognition, are freer to observe the data, and are freer to learn [14]. For an example about the usage of unsupervised learning to learn and reveal see [15].

2.1 Mental models

In order for us humans to make best use of the redundant sensory observations that we collect to build up our own mental models [16], our cognitive systems must make those models accessible to our future everyday perception. Mental models are perceived as internal representations of the environment that provide a conceptual framework for describing, explaining, and predicting future system states. These models should be “simple” (the parsimony principle for mental models, for example [17]), so that we will be able to use them to efficiently and quickly detect any new associations involved in learning. The mainstream cognitive psychology literature stresses that our brain is doing a profoundly difficult job in doing so [16]. It should deduce probabilistic links about our world and detect suspicious outliers, all by accessing and linking prior knowledge structures and schemas [18]. In this paper, we present a case study through which the process of creating the unsupervised ML models, as well as the models themselves are used to form, externalise and then articulate knowledge and, via that, making the learning more effective [19,20].

An organisation is typically a complex entity with many communicative channels connecting between the learning processes of its individuals and its whole culture [21]. A ML model, whether developed through supervised or unsupervised methods, will always be a simplification from a particular point of view on this complexity. This simplification and loss of detail is also a strength that enables new insights; and even more so when the “point of view” on the complexity is less determined by prior expectations, such as occurs with unsupervised methods.

The word ‘model’ is used in this paper to refer to different concepts interchangeably (sometimes deliberately). A ‘mental model’ is “produced through cognition by individuals to create a representation or structure of a phenomenon or solution to a problem” [23,16]. ‘To model’ is a verb describing the process that individuals undertake when they create, or retrieve existing mental models in order to solve problems [23]. A ‘computational model’ in this paper refers to unsupervised ML models, and the ‘co-learning’ process refers to a collaborative and connectivist process of developing the computational model.

Mental models help us explain and predict how learners interact with the world, and how they explain, understand, solve anticipated events, and communicate (see Figure 1 below from [23]). While mental models are internal structures [16], they can be exteriorised [24] when triggered by interaction with a domain system [25] such as robotics [23]. Our suggestion here is that unsupervised computational models can also offer an externalisation trigger. This allows for an observable effect of the initial internalisation of a mental model by the learners. By that, it can serve as a learning affordance, and the learning outcome can be observed through a change in learners’ language (for example by used concepts), and their ability to explain, predict and diagnose, as emphasised in [23] ‘transitory mental model’ (see the Figure 1 below), and as we will show in the results section.




Fig. 1. Mental model mode: diagram of functionality [23]

One of the atomic units of mental models are concepts, along with their role in language and communication [17]. Effective learning is strongly associated with developing a clear definition of concepts, including the meaningful relations between them [26]. Human reasoning and decision making are further based on this initial storage of concepts, which is one kind of mental representation. In an effective process of learning, the mental model will be stored [27] in the long-term memory of an individual, serving later as a schema [28], or a script [29]. Once the model has been created, it exists independently of its sources [30]. Visualisations, images and text can serve as mental affordances [31] or as we term them – learning affordances – by assisting with the functionality of short-term memory [32], reducing cognitive load, and therefore assisting learning. Our proposition is that unsupervised ML models can do that too, for example, by simplifying and reducing the number of the used dimensions.

One characteristic of mental models is that they are not immutable entities that remain invariant across (or even within) students [23]. Since they are subjective, they can lead to misconceptions [33,34]. They are channeled and processed through human long-term memory, which is essentially faulty, and thus they are not immune to biases and changes over time. Computational models, on the other hand, are more robust in that sense, as they are anchored on observed evidence and are externalised by statistical statements. In their work defining ‘mental affordances’, [31] specifies the three criteria for ‘something’ to be a mental affordance: (1) mental affordances are opportunities for mental action; (2) mental affordances are perceptible; and (3) the perception of a mental affordance involves the potentiation of the mental action that is afforded [31]. We raise the suggestion that unsupervised ML models can serve as a special kind of a mental affordance: a learning affordance. In Section 3 below we show how an unsupervised computational model serves as an opportunity for learning, how it is perceptible and that its perception involves the potentiation of learning.

2.2 Some criticisms of ML

Criticisms are continually raised about the role of ML in the context of its focus on prediction rather than explanation [35]. Many have argued that the use of ML in the learning sciences often lacks a theoretical basis [36, 37] and that it is often undertaken with “one eye closed to the peculiarities of the data” [38]. The latter is specifically challenging in the learning sciences domain, as many central constructs (such as

‘learning’ and ‘collaboration’) are often ambiguous in their operationalisation, which makes it a hard pre-requisite for supervised ML algorithm to work effectively. Many scholars have strongly argued that modeling should start with certain educational goals in mind (e.g., student learning outcomes), that are measurable [39, 40]. Here we show how the unsupervised approach help learning in a more open-ended manner by helping a learning organisation to make tacit into explicit and to adjust their long-held beliefs.

3 The co-learning process: Learning about trading behavior

The next subsections list the stages in which the unsupervised model was produced within our case study. This is essentially a very typical data mining process, but it illustrates the various ways in which the organisational co-learning process serves as a learning affordance. In particular, how it opens opportunities for mental action in a tangible way [31]. In being a ‘co-learning’ process, where the company’s expert trainers were engaged as designing partners and experts of their domain, each phase afforded them a different interaction with the data, and subsequently - learning opportunities, as described below. Figure 2 below summarises the main stages, showing where human semantics (of the experts) is playing an active role, when it is being guided by the data, and how new semantics is produced at the end of the pipeline. This dialogical interchange confronts the two (potentially differently biased) sources of knowledge, and by that potentially minimise the risks of a biased decision making (which by itself affords a new learning opportunity). Nevertheless, other mechanisms (out of the scope of this paper) such as ethical auditing, should also take place to minimise this risk further.

During the first phases, inputs such as experience and data guide the process, while the last two phases are unsupervised. The result is a parsimonious model, which affords new learning opportunities, that would potentially shape the mental model of the organisation (i.e., act as a learning affordance).

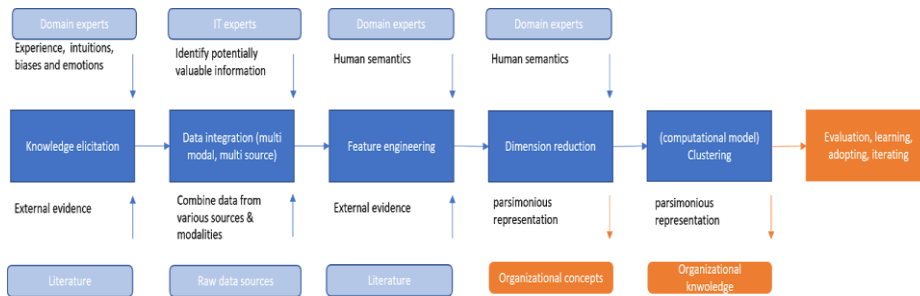


Fig. 2. A typical data mining process, resulting in an unsupervised model. The Blue rectangles represent the phases of the computing process. The light blue rectangles are semantic inputs going into the process. The Orange rectangles are outputs that go back to the organisation and potentially act as learning affordances.

Knowledge elicitation. During the knowledge elicitation phase, we triangulated evidence from the behavioral finance literature (to use empirical evidence about potential indicators for successful trading behaviors) with the domain experts’ tacit

semantics. This has ‘supervised’ our feature engineering and resulted in the domain experts’ externalisation of their existing mental model of the traders types (see Table 2 below).

Data integration. The complexity, inconsistencies and time associated with problem solving typically grow with the number of modalities or sources of information that need to be considered [41]. The resulting integrated data, while revealing an integrated semantics, is certainly one way in which computational models can be used as an affordance. In our case study, we used data from several data sources, such as clickstream from a trading software, various market characteristics, and internal reports for 2017-2018. As technically challenging and subject to the subjectivity of data sources the data integration phase is, its serendipitous nature opens new opportunities for learning.

Feature engineering. A typical feature engineering process is yet another opportunity for the human semantics to inflict itself and ‘supervise’ the modeling, while still allowing the set of statistical relationships to orchestrate a model without obvious human interference. At the end of this phase, we were left with a list of 35 features of traders’ behavior. Some features were as simple as the number of specific types of action that a trader makes in a typical month, such as the number of amendments and cancellations of orders. Other features were more ‘engineered’, such as the proportion of their trading in larger markets, and the proportion of the market they covered.

Dimension reduction. The parsimonious representation of a mental model makes it more accessible to short term memory retrieval, while also making it more vulnerable to flawed heuristics [17]. Dimension reduction is one way in which unsupervised ML offers some economy, and therefore removes some of the cognitive load related to high-dimensional models. From the computational side as well, the ‘curse of high dimensionality’ is a term used to describe the challenges introduced by the presence of a vast number of variables, resulting in a performance degradation [42]. In our case study, we used principal component analysis (PCA) on our 35 behavioral features and 480 traders’ data, after log-transforming the variables to make the distributions near-normal, and making sure that none of the variables were highly correlated with others. The PCA resulted in four main factors accounting for 83.37% of the total variance, which were then reviewed with the domain experts. The strongest factor was shown to cover elements of how active the trader is (quantifying the number of actions they initiated), the second was related to trading style (specifically, whether traders show the tendency to focus narrowly on a small number of markets, or whether they tend to widen their focus to cover a large number of markets, with less focus on each). The third factor covered the types of products they were typically traded (in terms of a product’s properties such as liquidity and volatility), and the fourth factor was focused on the overall volume of their trading. As will be shown in section 4, the introduction of these four dimensions into the organisational culture offered parsimony and a higher level of abstraction in the evolved mental model of trading behavior.

Clustering. Cluster analysis is a technique used to identify naturally occurring groups in a dataset, here used to profile trading behavior. In our case, the clustering was carried

out to challenge the existing profiling mental model of traders, that the company has been using in order to design a tailored training and mentoring programs. We deliberately did not add to the clustered features any feature having a direct relationship to performance measures (such as profit). The purpose was to make salient data patterns that were purely behavioral, to support formative feedback from mentors. The technical details of the carried K-means cluster analysis are detailed in the appendix.

4 Results: mental models of trading behavior

To validate our hypothesis that the unsupervised model had affected the mental model of the training managers in the organisation, we interviewed two of the most senior ones. The interview questions, as well as the themes we followed in the analysis were based on the ‘transitory mental model’ [23] dimensions, such as the models’ effects on language, prediction, diagnosis and supporting learners. Below are the main findings. In the interview the managers were asked to compare the mental model created within the company prior to our analysis, with the computational model created by that analysis (see Tables 2 and 3 respectively in the appendix).

The main differences between the two models: The main difference noted was that the first model was “*subjective*”, in the sense that it had been derived from long experience of being traders and trainers, whereas the second model was “*objective*”, in the sense that it had emerged from the trading data. A related difference was in the number of profiles. The managers felt that they could have composed their mental model with more than the five profiles that they did, but chose not to so, as they had no easy way to determine what would be a sufficient set to cover the field. By contrast, arriving at four profiles rather than some other number was driven by the usual needs for parsimony vs. coverage of the data in unsupervised ML. Another important difference was that the second model more clearly articulated “*how engaged a trader is*” compared to the first model as it brought to the fore issues around order activity and diversity. This highlighted the fact that there were few tools available to managers to measure engagement with the data available to them.

The main commonalities between the two models: One manager suggested that “*both models very much focus on the markets [and] the types of markets, and the characteristics of the markets in terms of what the individuals seem to prefer and pay particular attention to volatility and product diversity*”. An indirect issue is that neither model in itself indicates how best to train a particular type of trader, but the insight that training and mentoring could, in principle, be adapted in such a way indicated a shift in perception.

The ML model’s effect on the managers’ views on trading and training: On the positive side, it was noted that the ML model was able to show up similarities and differences between mixes of trader behavior in the various company offices across the world. It was also noted that the co-learning process had demonstrated the value of the company’s data. One manager noted the value of the ML model in supporting both recruitment and mentoring, in a way that was not really possible with the initial model. In terms of recruiting, the ML model opened up the possibility of consciously improving the diversity of trading styles within the company as part of its overall risk management strategy. On the negative side, “*One disadvantage would be the potential*

for assumptions and first impressions to be derived from cluster categorisation before a mentor has really . . . had a chance to work closely with and get to know a trader” and might also have a similar downside in recruitment. An interesting observation was made about the possibility that the ML profiles might become too rigid and constraining, and the need to update it when markets change.

The ML model’s effect on the managers’s language: Several ideas emerged with regard to language and concepts. From one manager, *“the key thing for me . . . was about behaviors, and how traders’ behaviors [our emphasis] with the markets”*. This contrasts with the former focus on performance, typically profit and loss. For the other manager, *“the most obvious change is recognition that traders can be grouped by certain factors, other than their start date, and that they don’t have to be considered as individual entities at all time[s].”* This manager also pointed out the consequence of the objectivity of the data, *“Due to the mentality of traders and analysts, and generally all departments within ZISHI/OSTC, I think there is a stronger propensity for people to acknowledge and factor in advice when it originates from a data led approach.”* The ML model has resulted in greater clarity, *“Both models, but certainly the newer one, separate volume and position size as separate factors. I think evidencing this is important as a common incorrect assumption is that a bigger position equates to higher volume. This teamed with the concept of complexity really broadens the criteria on which a traders’ activity is ‘judged’.”* Note that the term “complexity” itself, which is now used routinely within the company, derived from the unsupervised modeling process.

The ML model also helped to extend the manager’s thinking about the evolution of a trader’s behavior with experience, *“The concept and likelihood of a trader gravitating from one cluster to another over time, which has been mentioned in the analysis, but perhaps not in the new model cluster descriptions is one aspect I think mentors and decision makers internally either need to accept as a possibility or on the other scale accept isn’t as straightforward as they assume.”*

The ML model’s effect on the managers’ ability to ‘diagnose’ or ‘predict’: One manager anticipated that, *“having a data fed model gives a mentor more confidence to forward plan, anticipate and react quicker to the obstacles and barriers that each trader may face when progressing. It also adds greater weight to your case when presenting why you think a trader needs more time/resources to develop than a higher-level decision-making body may invoke. It also has the ability to make communicating and justifying a mentor’s approach and ‘diagnosis’ to others easier. This could even be as simple as having the reassurance that the same language and concepts are being used.”*

Are there any other changes, not already noted, the managers have noticed in themselves and in their colleagues: One manager explained that, *“This type of insight, . . . does open our eyes to who we have as a workforce and who traders are and how you can really look at performance.”* The other pointed out that: *“there is a growing recognition that pairing the right mentor and mentee could be an important part of getting it right – although still somewhat in its infancy. From my own perspective I’m certainly more conscious of the fact that certain characteristics, or even whole clusters, that are considered to result in less profitable performance within ZISHI/OSTC might actually suit alternative trading environments which work under different parameters. . . . I certainly think it’s made me consider if dependent on cluster characteristics, and*

very much linked to this is product mix, if risk parameters need to be adapted for the different clusters, even at a very junior stage of a career.”

In summary, the main affordance for the training managers and trainees is that the profiles derived by unsupervised ML created a handy, bias-fencing shorthand to encapsulate a large number of low-level behavioral variables. These behavioral variables are usually not directly observable by the managers themselves before the modelling, and developing a mental model would typically take significant cognitive effort and time. In addition, the fact that those models were unsupervised has afforded the trainers a sense of validation, as well as of standardisation across the company’s different international locations. The unsupervised model and the concepts arising from the modeling changed the language within the company. New concepts such as “*complexity*”, “*cluster*”, and “*trading style*” were introduced and diffused across the company. Lastly, the new model opened up new dialogues about the shift of focus from performance onto behavior, about helping traders mobilise between profiles as part of their progress. It also opened up the possibility for more targeted recruitment and mentoring, as well as for potentially better matching traders to types of markets or mentors. The managers also identified a potential risk of trainers using the clusters as too simplified or rigid, and mentioned the need for awareness, and re-modeling.

5 Conclusions

The main rationale for using unsupervised ML models is that they can expose unexpected patterns, and therefore adds data-driven semantics to the existing semantics of human experts. Of course, that does not come without challenges and dangers. In this paper, we have used a case study of a trading and training company managers which reflected on the comparison between their own mental model, and a computational model produced by a co-learning process. We suggest that unsupervised models can be referred to as learning affordances, as they have the potential of reducing the complexity of a highly-dimensional behavior, floating inexplicit or unexpected patterns, introducing new concepts to the company’s language and generally affecting its learning. For example, the company’s training managers were now able to discuss how traders can be encouraged to move between different clusters of behavior and how their behavior might relate to performance. The model was an outcome of a collaborative modeling process, which suggests that it was itself a learning affordance. In our case it has enabled a space and time for a continuous interaction between the experts’ tacit knowledge, their own various mental models and data modeling. The method used in this case study should be generalisable to other unsupervised methods and to organisations that have accumulated a large amount of untapped behavioral data.

Appendix A: Technical details of the cluster analysis

Can be found at [Technical details](#).

References

1. Seel, N. M.: Model-based learning: a synthesis of theory and research. *Educational technology research and development*, 65(4), 931–966. doi:<https://doi.org/10.1007/s11423-016-9507-9> (2017).
2. du Boulay, B., O'Shea, T., & Monk, J.: The black box inside the glass box: presenting computing concepts to novices. *International Journal of Man-Machine Studies*, 14(3), 237-249. (1981).
3. Marsick, V. J., & Watkins, K. E.: Demonstrating the value of an organisation's learning culture: the dimensions of the learning organisation questionnaire. *Advances in developing human resources*, 5(2), 132-151. (2003).
4. Echeverria, V., Martinez-Maldonado, R., & Buckingham Shum, S.: Towards collaboration translucence: Giving meaning to multimodal group data. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (pp. 1-16). (2019).
5. Kent, C., Chaudhry, M.A., Cukurova, M., Bashir, I., Pickard, H., Jenkins, C., du Boulay, B., and Luckin, R. (2021). On how unsupervised learning can shape minds: a very brief overview. In *Proceedings of the 11th international conference on learning analytics and knowledge* (pp. xx).
6. Ambrose, S., Bridges, M., DiPietro, M., Lovett, M., & Norman, M.: *How learning works: 7 Research-Based principles for smart teaching*. San Francisco: Jossey-Bass. (2010).
7. Mayer, R. E.: *The promise of educational psychology, volume 2: Teaching for meaningful learning*. Upper Saddle River, NJ: Merrill Prentice Hall. <https://onlinelibrary.wiley.com/doi/abs/10.1002/pfi.4930420410>. (2002).
8. Watkins, K. E., & Kim, K.: Current status and promising directions for research on the learning organisation. *Human Resource Development Quarterly*, 29(1), 15-29. (2018).
9. Nonaka, I., & Takeuchi, H.: *The knowledge creating company*. Oxford, England: Oxford University Press. (1995).
10. Siemens, G.: Connectivism: A learning theory for the digital age. *International Journal of Instructional Technology and Distance Learning (IRRODL)*, 2(1), 3–10. (2005).
11. Smart, P., Engelbrecht, P., Braines, D., Hendler, J., & Shadbolt, N.: The Extended Mind and Network-Enabled Cognition. *The Computer Journal*. (2008).
12. Fjortoft, H., and Lai, M. K.: Affordances of narrative and numerical data: A social-semiotic approach to data use. *Studies in Educational Evaluation*, 100846. (2020).
13. Gibson, J. J.: *The ecological approach to visual perception: classic edition*. Psychology Press. (2014).
14. Amershi, S., and Conati, C.: Combining unsupervised and supervised classification to build user models for exploratory learning environments. *JEDM| Journal of Educational Data Mining*, 1(1), 18-71. (2009).
15. Zhang, N., Biswas, G., and Dong, Y.: Characterising students' learning behaviors using unsupervised learning methods. In *International Conference on Artificial Intelligence in Education* (pp. 430-441). Springer, Cham. (2017).
16. Johnson-Laird, P. N.: *Mental models: Towards a cognitive science of language, inference, and consciousness* (No. 6). Harvard University Press (1983).
17. Holyoak, K. J., and Morrison, R. G. (Eds.): *The Cambridge handbook of thinking and reasoning* (Vol. 137). Cambridge: Cambridge University Press. (2005).
18. Barlow, H. B.: Unsupervised learning. *Neural computation*, 1(3), 295-311. (1989).
19. Bransford, J. D., Brown, A. L., and Cocking, R. R.: *How people learn: Brain, mind, experience, and school* (Expanded edition). Washington, DC: The National Academies Press. (2000).
20. Sawyer, R. K. (Ed.): *The Cambridge handbook of the learning sciences*. Cambridge University Press. (2005).

21. Marsick, V.J. and Watkins, K.E.: *Facilitating Learning Organisations*, Gower, Brookfield, VT. (1999).
22. Fincher, S., Jeuring, J., Miller, C. S., Donaldson, P., du Boulay, B., Hauswirth, M., . . . Petersen, A.: *Capturing and Characterising Notional Machines*. Paper presented at the ITiCSE '20: Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education, Trondheim, Norway. (2020).
23. Edwards-Leis, C. E.: *Challenging learning journeys in the classroom: Using mental model theory to inform how pupils think when they are generating solutions* (2012).
24. Barker, P., van Schaik, P., Hudson, S., and Meng Tan, C. 1998. *Mental models and their role in the teaching and learning of human-computer interaction*, Volume 1 of Proceedings of ED-MEDIA/ED-TELECOM 98, 10th World Conference on Educational Multimedia and Hypermedia, Carroll, J. M., and Olson, J. R. 1988. *Mental models in human-computer interaction*. In *Handbook of human-computer interaction*. Elsevier. 45–65. Freiburg, Germany, June 20-25, T. Ottman and I. Tomek (Eds.), Association for the Advancement of Computing in Education, Charlottesville, Virginia, USA. (1998).
25. Carroll, J. It. and Olson, J.R.: *Mental Models in Human-Computer Interaction*, in *Handbook of Human-Computer Interaction*, M. Helander (ed.), Elsevier Science Publishers, Amsterdam, pp. 45-65. (1988).
26. Ausubel, D. P.: *Educational psychology: A cognitive view*. New York: Holt, Rinehart & Winston. (1968).
27. Bucciarelli, M.: *How the construction of mental models improves learning*. *Mind & Society*, 6(1), 67-89. (2007).
28. Anderson, R. C.: *The notion of schemata and the educational enterprise: General discussion of the conference*. In R. C Anderson, R. J. Spiro & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge*. (1984). Hillsdale, NJ: Lawrence Erlbaum. (1977).
29. Preece, J., Rogers, Y., Sharp, H., Benyon, D., Holland, S., and Carey, T.: *Human computer interaction*, NJ: Addison Wesley. (1994).
30. Barker, P. G.: *Mental models and network pedagogy*, Conference proceedings of EN-ABLE99, International Conference EVITech, Helsinki University, Finland. (1999).
31. McClelland, T.: *The mental affordance hypothesis*. *Mind*, 129(514), 401-427. (2020).
32. Henderson, L. and Tallman, J.: *Stimulated recall and mental models*. Lanham, ML: Scarecrow Press, Inc. (2006).
33. Vosniadou, S., and Brewer, W.F.: *Mental models of the earth: A study of conceptual change in childhood*. *Cognitive Psychology*, 24(4), 535–585. doi.org/10.1016/0010-0285(92)90018-W (1992).
34. Franco, C., and Colinviaux, D.: *Grasping mental models*. In *Developing models in science education* (pp. 93-118). Springer, Dordrecht (2000).
35. Vaughan, J. W., and Wallach, H.: *A Human-Centered Agenda for Intelligible Machine Learning*. *Machines We Trust: Getting Along with Artificial Intelligence*. (2020).
36. Nelson, L. K.: *Computational grounded theory: a methodological framework*. *Sociol. Methods Res.* 49:0049124117729703. doi: 10.1177/0049124117729703 (2017).
37. Penuel, W. R., and Shepard, L. A.: *Assessment and teaching*. In D. H. Gitomer, & C. A. Bell (Eds.). *Handbook of research on teaching* (pp. 787–850). (5th ed.). Washington, DC: American Educational Research Association. (2016).
38. Radford, J., and Joseph, K.: *Theory In, Theory Out: The uses of social theory in machine learning for social science*. arXiv, arXiv-2001 (2020).
39. Hamilton, L., Halverson, R., Jackson, S., Mandinach, E., Supovitz, J., and Wayman, J.: *Using student achievement data to support instructional decision making* (NCEE 2009-4067). Retrieved from Washington, DC: National Center for Education Evaluation and Regional Assistance, Institute of Education Sciences, U.S. Department of Education. <http://ies.ed.gov/ncee/wwc/publications/practiceguides/> (2009).

40. Mandinach, E. B., and Schildkamp, K.: Misconceptions about data-based decision making in education: An exploration of the literature. *Studies in Educational Evaluation*, 100842. (2020).
41. Rapp, D. N.: Mental models: Theoretical issues for visualisations in science education. In *Visualisation in science education*. pp. 43-60. Springer, Dordrecht. (2005).
42. Han, J., Pei, J., and Kamber, M.: *Data mining: concepts and techniques*. Elsevier. (2011).
43. Kaiser, H. F.: An index of factorial simplicity. *Psychometrika*, 39, 32-36 (1974).
44. Aldenderfer, M. S., and R. K. Blashfield: *Cluster Analysis*. Newbury Park: Sage Publications (1984).