Investigating LC3 and influenza M2 trafficking via robust representation learning on high throughput microscopy datasets

Omar El Oakley, 31st of August 2022

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Philosophy (MPhil, Division of Biosciences) at University College London (UCL)

I, Omar Reina El Oakley, 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 thesis. Signed,

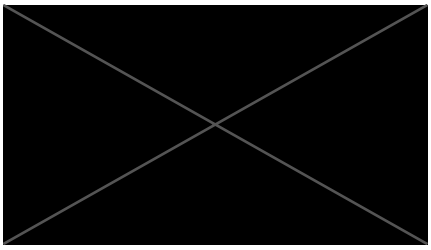

To my mother,

Gillian El Oakley

Acknowledgements

I thank Dr R. Beale for his experimental guidance and scientific advice, Dr M. Johnson for his invaluable insights into deep learning and good coding practices (all extant mistakes in the code base are, of course, my own), Dr E. Marcassa and J. Li for their help getting set up with the ImageStream assay (and the occasional coffee break), S. Isaksson for his help with performing image data normalisation and the quirks of training GANs, Dr R. Ulferts for her support and advice, Dr B. Montaner for her help with cell culture and molecular biology techniques, Dr E. Carr for his discerning eye for the proper plotting of statistical figures, and finally L. Timimi and C. Figueras-Novoa for their insights, training, and experimental and emotional support.

I thank the members of my thesis committee, Dr S. Tooze, Dr K. Bentley, and Dr R. Henriques for their support, encouragement and guidance.

Abstract

Quantitative descriptors of bioimage features have been used to efficiently discriminate between different phenotypic populations within microscopy image datasets. However, trying to discriminate between the processes of canonical and non-canonical autophagy is difficult as they share a key patterning determinant - relocalisation of LC3 - and because both responses are subject to significant heterogeneity. Here, we apply both classical feature extraction and unsupervised representation learning to isolate and the describe the differences between canonical and non-canonical autophagy. In doing so, this thesis aims to demonstrate how using quantitative bioimage analysis in this context can increase the robustness of observations derived from hypothesis-driven high throughput imaging experiments. A texture descriptor

approach is contrasted with two deep representation learning methods, the InfoGAN and BYOL, where it is seen that classical approaches are more amenable to statistical testing and allow for stronger inferences about the underlying biology. These classical approaches are then used to identify imaging phenotypes unique to different stimuli of non-canonical autophagy. From a methodological perspective, this thesis underlines the continued utility of classically-derived image features for scientific discovery.

Impact Statement

Deep learning enabled image analysis is being widely investigated as a way to supplement and inform decisions in a wide range of fields, including medical imaging and robotics. Exploring how these networks undertake visual reasoning, and in particular the ways in which this visual reasoning can diverge from how expert practitioners view their subject matter, will be essential in ensuring these systems function safely and fairly. Additionally, understanding how we can better use quantitative information to supplement directly observational experiments will hopefully allow for more targeted formulations of quantitative hypotheses, reducing time spent on low-priority avenues of enquiry and in doing so reducing research waste.

Table of Contents

Abbreviations

BYOL - Bootstrap Your Own Latent
CASM - conjugation of atg-8 to single membranes
EGFP - enhanced green fluorescent protein
eIF2α - eukaryotic initiation factor 2 alpha
GAN - generative adversarial network
HA - influenza haemagglutinin protein
InfoGAN - information maximizing generative adversarial network
JSD - Jensen-Shannon divergence
LC3 - microtubule-associated protein 1 light-chain 3
LDA - linear discriminant analysis
LIR - LC3-interacting region
3-MA - 3-methyladenine
mTORC1 - mammalian target of rapamycin complex 1
NS1 - influenza non-structural protein 1
PB2 - influenza protein polymerase basic 2
PI3K - phosphoinositide 3-kinase
PI3P - phosphatidylinositol 3-phosphate
PKR - protein kinase R
SIE - superinfection exclusion
TGN - *trans*-Golgi network
TRIM - tripartite motif-containing protein
ULK1 - *unc*-51 like kinase 1
vATPase - vacuolar ATPase
VPS34 - *vacuolar protein sorting* 34

1 - Introduction and Aims

In autophagy research, the lipidation of a core autophagy marker, LC3, is used as a semi-quantitative readout for the process of canonical autophagy. This lipidation process drives a relocalisation of LC3 from cytoplasmic to membrane compartments, resulting in a distinct imaging phenotype, where diffuse LC3 is collected into discrete puncta. However, LC3 lipidation can take place in contexts functionally and molecularly dissimilar to canonical autophagic pathways, contexts sometimes referred to as non-canonical autophagy (Table 1). Reasoning directly from LC3 lipidation to canonical autophagy has led to contradictory interpretations and confused inferences in contexts where non-canonical autophagy is seen to occur. Difficulty in measuring and defining phenotypes in this context is compounded by heterogeneity inherent to both processes. This thesis will outline a set of methods that distinguish canonical from non-canonical autophagy when imaging LC3, and subsequently make inferences about cellular behaviour under non-canonical autophagy, inferences which may have otherwise been ill-constructed if premised on naive interpretations of LC3 lipidation.

	Canonical autophagy	*Non-canonical autophagy*
Site of lipidation	Double membranes	Single membranes
Stimuli	Nitrogen starvation Dysfunctional organelles Protein aggregates Cytosolic pathogens	Ionophore drugs LC3-associated phagocytosis Aberrant ion channel activity Bacterial secretion systems
Function	Lysosomal degradation Metabolic homeostasis Organelle maintenance Microbial restriction	Unknown
Molecular Determinants	ULK1 complex VPS34 complex Atg9/Atg2	vATPase complex ATG16L1 WD40 domain

Table 1: An outline of key differences between canonical and non-canonical autophagic processes.

1.1 The Influenza A M2 Protein

Influenza A is a respiratory virus that is capable of causing pandemics. It possesses a segmented, negative-sense RNA genome, and aquatic birds are the primary natural reservoir for the virus, prompting frequent viral crossover from domestic poultry populations (Parrish, Murcia, and Holmes 2015). Understanding the mechanisms by which it avoids and subverts cellular defenses is essential to reduce both the scale of its impact on global health and the possibility of emergence of highly transmissible and pathogenic viral strains (Pappas et al. 2008; Lam et al. 2011). Due to the highly restricted coding capacity of the viral genome (Dubois, Terrier, and Rosa-Calatrava 2014), most influenza proteins perform multiple roles during the

viral lifecycle. The influenza M2 protein acts primarily as an ion channel, and its best described role is to acidify the virion upon endosomal entry, triggering release of the viral genome from virion into the cytoplasm (Shaw and Palese 2007). Additionally, the haemagglutinin (HA) cell entry complex is acid-sensitive and must be prevented from triggering during transport through the acidic trans-Golgi network (TGN). As such, the M2 protein has a parallel role during viral exit where it acts to neutralise ion gradients across a set of secretory endomembrane compartments (Ciampor et al. 1992; Alvarado-Facundo et al. 2015). A further proposed role for the M2 protein can also be found during the budding of new virions from the plasma membrane, where structural studies have shown a contribution of the transmembrane portion of the protein to membrane curvature and scission (reviewed in Rossman and Lamb 2011).

Despite the experimental focus on the ion channel activity of the M2 protein, M2 additionally possesses the longest cytoplasmic tail of any influenza A protein (Iwatsuki-Horimoto et al. 2006), with a complex set of post-translational modifications and proposed interaction partners (Holsinger et al. 1995; Tripathi et al. 2015). As there is such a strong coding capacity constraint, the presence of this cytoplasmic tail is strongly suggestive of M2 possessing functions beyond endomembrane ion gradient disruption, possibly as a signalling interface to the host cell. Previous work in our laboratory has defined and explored one such interaction, that of M2 with members of the Atg8 protein family (typically referred to, in aggregate, as ‘LC3’ in mammalian contexts) via a highly conserved LC3-interacting region (LIR) motif. This protein is a universal component of the autophagy machinery in eukaryotes.

1.2 Autophagy and the Influenza A Virus

1.2.1 Macroautophagy and its regulators

Autophagy refers to a set of processes in which bulk cytoplasmic components are degraded via delivery to the lysosome. The term is first defined in the context of the discovery of the lysosome (de Duve et al. 1955), in deliberate contrast with the process of delivery of *extracellular* contents to the lysosome (then defined as heterophagy, now more commonly known as endocytosis). Autophagic processes are critical under conditions of low nutrient availability as the products of lysosomal degradation can be recycled as substrates for cellular respiration and anabolic reactions. This criticality of autophagy to the cellular starvation response underpinned the yeast genetic screens that delineated the molecular determinants of autophagy - a group of thirty proteins that are almost universally conserved in eukaryotes (reviewed in Ohsumi 2014). Subsequent research found that autophagy functioned outside of starvation by selectively delivering dysfunctional or damaging cellular components, such as aging mitochondria or large protein aggregates, to the lysosome for destruction - thereby contributing to cellular maintenance (reviewed in Johansen and Lamark 2011). Outside of cellular maintenance, the capacity of autophagy to destroy contents of the cytoplasm can be utilised by the cell as a restriction mechanism for a variety of intracellular pathogens, in a set of processes described as xenophagy (reviewed in Wileman 2013). The growing understanding of xenophagy has led to a variety of descriptions of bacterially and virally-encoded systems of disruption or subversion of autophagic machinery (reviewed in McEwan 2017).

Macroautophagy is a form of autophagy which generates distinctive double membrane structures known as phagophores to engulf cellular contents prior to their delivery to the lysosome (summarised in Ravikumar et al. 2009). This requires finely tuned orchestration of cellular signal transduction, substrate recognition, membrane trafficking and remodelling. To

parse this complexity, a protein known as LC3 is often used to interrogate autophagic processes, firstly because the lipidated form of the protein is specifically and strongly localised to autophagic membranes and thus allows profiling of autophagic structures with fluorescence microscopy. Secondly, because processed and unprocessed populations of LC3 can be easily resolved via western blot, and the relative proportion of these two populations can be informative about the rate of flux through the pathway (Klionsky et al. 2021).

Two kinase complexes cooperate to integrate cellular nutrient sigalling and induce the generation of autophagic membranes, and thus occupy central roles in the literature. The ULK1 complex is a protein kinase complex that is directly inhibited by mTORC1 binding. In starvation conditions, mTORC1 dissociates, releasing the inhibition of the protein kinase activity of the complex. While it is likely ULK1 possesses a large number of protein substrates, two key substrates, Vps34 and Beclin1, are found in the other central autophagic kinase - the class III phosphoinositide 3-kinase (PI3K) complex. Phosphorylation of Vps34 and Beclin1 in turn drives the phosphatidylinositol 3-phosphate (PI3P) kinase activity of this complex. After localisation to a nascent autophagic membrane, these two kinase complexes act in concert to create an autophagosome specific pool of PI3P. Molecular recognition of the PI3P pool allows for the recruitment of the LC3 lipidation machinery, consisting at its core of the key complex ATG12-ATG5-ATG16L1, which acts as an E3-like ligase by conjugating the ubiquitin-like LC3 to lipid molecules, most often phosphatidylethanolamine. As these kinases sit at the very top of the autophagy signalling cascade and are thus essential for the generation of autophagosomes, they often undergo genetic or chemical manipulation as a method of investigating the contribution of autophagy to other cellular processes (reviewed in Lamb, Yoshimori, and Tooze 2013).

1.2.2 Influenza infection and canonical macroautophagy

The influenza virus non-structural protein 1 (NS1) possesses a well described protein kinase R (PKR) inhibitory activity. This acts to limit activation of the interferon signalling system and thus constrains the activity of the innate immune system (Bergmann et al. 2000). As PKR also sits upstream from eukaryotic initiation factor 2α (eIF2α), the early literature concerning the relationship between autophagy and influenza posits it is primarily antagonistic with virally-derived autophagic *inhibition* functioning to reduce innate immune activity and viral antigen presentation (Randall and Goodbourn 2008; Schmid and Münz 2005; Orvedahl and Levine 2008). The first paper describing a direct measurement of autophagic activity in influenza A infection demonstrated a significantly increased degree of LC3 lipidation via SDS-PAGE (Zhou et al. 2009). It was additionally seen that a pharmacological inhibitor of autophagy, 3-methyladenine (3-MA) seemed to reduce viral titres, concluding from this that influenza induced autophagy and this induction directly contributed to viral fitness - in contrast to predictions about autophagy acting as host defense. However, the validity of 3-MA as an autophagic inhibitor has been challenged (Y.-T. Wu et al. 2010), in that it has been shown that it can in fact induce or potentiate autophagy in nutrient-rich mammalian contexts.

Gannagé et al. demonstrated a similar increase in lipidated LC3 under infection, but showed that there was no further increase in lipidated LC3 under inhibition of lysosomal flux. This was interpreted as influenza acting to inhibit the final steps of the macroautophagy pathway, where the contents of autophagosomes were prevented from being degraded in lysosomes (Gannagé et al. 2009). It was further found that the influenza M2 protein alone could induce a similar autophagic phenotype under transient transfection, and that the identified autophagosomes induced under these conditions had radically differing morphologies and transport kinetics to those induced under conditions of canonical macroautophagy when visualised under confocal

fluorescence microscopy. The paper describes a further set of experiments which show that the M2 protein induces this autophagic blockade independently of its ion channel activity. Amantadine was employed to specifically block the M2 ion channel, and it was seen that upon amantadine treatment there was no impact on LC3 lipidation. Notably, these experiments were conducted in a strain known to be resistant to amantadine treatment (Appleyard 1977).

In contrast to the findings of Gannagé et al., a closer investigation of the role of the cytoplasmic tail in the potentiation of increased LC3 lipidation showed a partial dependence on the LIR motif (Beale et al. 2014), in contrast to the localisation of function to the N-terminal region. Additionally, it was shown that the increased LC3 lipidation was independent of the FIP200 complex - a component of the ULK1 complex and thus key upstream regulator of canonical autophagy - while being dependent on the core lipidation machinery through ATG16L1. A further contribution of Beale et al. was in describing a distinctive LC3 patterning phenotype, in which the LC3 molecule localised strongly to the plasma membrane - a unique phenotype not seen in canonical autophagy. Secondly, the independence of the LC3 lipidation phenotype from M2 ion channel activity was explicitly overturned (Ren et al. 2016; Fletcher et al. 2018), via the use of amantadine-sensitive influenza strains.

Later papers which directly investigate the interaction of influenza with autophagy include (Comber et al. 2011) which attempts to directly investigate the effects of influenza-derived autophagy on antigen presentation, but relies on pharmacological tools with many off target effects in a similar manner to (Zhou et al. 2009). Interestingly, an assay dependent on high-throughput imaging flow cytometry demonstrated that LC3 relocalisation induced by influenza infection could be differentiated from canonical autophagy as it allowed for high LC3 relocalisation in addition to strong caspase-3 cleavage (de la Calle et al. 2011), a surprising

result as canonical autophagy is typically seen to be downregulated prior to initiation of apoptosis (Xu, Yuan, and Lipinski 2013).

Work in highly pathogenic H5N1 avian influenza suggested that virally-induced autophagy acts to increase viral pathogenicity (Ma et al. 2011), and in contrast to Gannagé et al., proposes that autophagy (and attendant “autophagic cell death”, which is in itself a contentious concept (Shen, Kepp, and Kroemer 2012; Kroemer and Levine 2008)) is induced rather than stalled at the autophagosome-lysosome fusion stage. However, there are limited experimental controls for cell entry or lowered viral replication phenotypes when assessing virus-induced cell death, limiting the strength of its conclusions. This line of reasoning linking high pathogenicity to autophagic signalling is continued in mouse models (Sun et al. 2012), however there is little recognition that increased LC3 accumulation may result from an inhibition of flux (despite (Gannagé et al. 2009)) and the reliance on 3-MA makes the paper’s strong causal conclusions about inhibiting autophagy (for example, via the use of chloroquine (Yan et al. 2013)) as a clinical intervention problematic. The (Gannagé et al. 2009) and (Zhou et al. 2009) papers result in two different patterns of citation, with multiple review papers (and research papers concerned with the interactions of autophagy with other viruses) diametrically claiming that influenza either inhibits or induces autophagy, respectively - resulting in a somewhat confused literature concerning the subject of autophagy in influenza.

Complicating the autophagy-pathogenicity relationship, experiments with H5N1 pseudotyped viral particles demonstrated strong induction of a form of autophagy which was dependent on the presence of haemagglutinin - as opposed to M2 - and the presence of upstream autophagic regulators such as Beclin-1 (Pan et al. 2014), suggesting there may be multiple pathways of autophagy induction present, dependent on the particular strain of influenza used and its pathogenicity. This is consistent with similar work exploring the effects of a host-restrictive

polymorphism in the influenza polymerase subunit PB2 (Kuo et al. 2017), finding that only avian-signature viral titres were significantly affected by a set of autophagic genetic manipulations. These papers clarify - to a certain extent - earlier contradictory work as they suggest a distinct pattern of autophagy induction specific to highly pathogenic avian influenza, with an effect on viral titre. This is in contrast to human- or lab-adapted influenza strains which - in cell culture - are unaffected by autophagic manipulation and are independent of canonical upstream regulators of autophagy. Later work using carefully constructed cellular models deficient in core autophagy lipidation machinery (Atg5) seems to confirm the lack of a concrete effect of autophagy on viral titre (Perot et al. 2018), demonstrating instead an autophagy-dependent reduction in interferon expression. There is some work, however, that puts forward evidence that Beclin-1 overexpression, a canonical upstream autophagy regulator, plays a titre-modulating role in H1N1 infection (Feizi et al. 2017), but the interferon-inducing effects of RNA transfection are not controlled for - an essential consideration in RNA virus research (X. L. Li et al. 1998).

Further work on the relationship between upstream autophagic signalling and influenza infection suggests that viral proteins, including M2 and HA, can act to induce increased mTORC signalling, although there is no investigation into the subsequent effects on autophagy induction (Kuss-Duerkop et al. 2017). This suggests that upstream regulators of autophagy are undergoing some form of viral manipulation, but are not the key points of control for viral autophagy. As the mTORC complexes sit upstream of a wide variety of virus-relevant processes - notably transcriptional regulation - it is likely that induction of mTORC signalling primarily acts to potentiate virus protein production. Additionally, work investigating knockdown of Rubicon, a key autophagy inhibitor, under influenza infection demonstrated a subsequent reduction in interferon signalling (Wan et al. 2017). However, as above, it is unclear whether this is mediated by Rubicon’s effects on autophagy.

Parallel to the work on viral replication dynamics, a well-characterised canonical autophagic haploinsufficiency mouse model (Yue et al. 2003) was used as a basis to explore the dendritic cell host response under canonical autophagy deficiency, demonstrating a partial dependence on functional canonical autophagy for full T-cell activation and antigen cross-presentation under influenza infection (Zang et al. 2016), corroborating to a certain extent the results found in (Comber et al. 2011).

Additional work compares the effects of deletions across a set tripartite-motif containing (TRIM) proteins, a set of interferon effectors with E3 ubiquitin-ligase activity, on autophagsome induction across different viruses. By manually counting autophagosome number in TRIM-deficient backgrounds, a TRIM23 axis was found to be a unifying signalling pathway linking interferon induction to autophagosome formation in viral infection, including H1N1 influenza infection (Sparrer et al. 2017). This is of particular methodological interest as the validity of this analysis is premised on the assumption that autophagy induction only manifested via the formation of individuated autophagosomes, despite the fact that influenza-induced autophagy was seen to have radically non-standard patterning under fluorescence microscopy (Gannagé et al. 2009; Beale et al. 2014), highlighting the difficulty of trying to perform comparative phenotypic analysis in settings where the phenotype itself is either non-standard or not easily defined.

In summary, the central question of whether autophagy is host-protective or is virally subverted is subject to confusing and contradictory claims. Influenza directly manipulates upstream regulators of autophagy, although possibly to effect changes in processes unrelated to autophagy. It possesses a highly conserved interaction motif with a core autophagy marker and yet (at least in strains adapted to human hosts) the complete abrogation of central autophagic

machinery has no effect on viral titre. When autophagy is allowed to be induced its primary effect seems to be to limit innate immune signalling, prompting the question as to why influenza would derive an increase in fitness from actively inhibiting a host response that limits interferon activity.

1.2.3 Influenza infection and non-canonical autophagy

Fletcher et al. demonstrate that much of the above interpretational confusion originates from the inference of canonical autophagy directly from LC3 lipidation data. A hallmark of canonical macroautophagy, as outlined above, is that LC3 is inserted into double membrane structures. A key result in the paper is that a non-canonical system of conjugation of LC3 to single membranes, prototypically described in LC3-associated phagocytosis (Florey et al. 2011), can be differentiated from canonical autophagy as it is specifically abolished via the removal of the C-terminal WD40 domain of the ATG16L1 protein (Fletcher et al. 2018). Without this domain, canonical autophagy can still proceed, yet no lipidation is seen under influenza infection, strongly suggesting that influenza-induced LC3 lipidation only occurs in this non-canonical context. This form of non-canonical autophagy, later described as conjugation of atg-8 to single membranes (CASM), is known to be independent of the nutrient-homeostasis signalling pathways that drive canonical autophagy. This is consistent with the independence of influenza-derived autophagy from upstream canonical autophagic signalling seen in some parts of the earlier literature. As a result, reasoning about the functional and evolutionary relationships of influenza to LC3 lipidation necessarily must be divorced from the known functions of canonical autophagy, and thus makes the analysis of previous results in the literature complex. Further description of this process has shown that it is dependent on the vacuolar ATPase (vATPase) complex as a key signal transducer, whereby a host of stimuli that induce aberrantly neutral ion gradients in the endolysosomal compartment (including ionophore drugs such as

monensin alongside the M2 protein) induce vATPase-mediated recruitment of the ATG16L1 complex via its WD40 domain (Ulferts et al. 2021; Hooper et al. 2022).

What, then, is the function of this non-canonical form of autophagy in influenza infection? Results from a mouse model lacking the ATG16L1 WD40 domain demonstrate a host-detrimental phenotype whereby low-pathogenicity strains become lethal in the absence of non-canonical autophagy (Yingxue Wang et al. 2021). Bone marrow transplant experiments with wild-type mice show that the phenotype is not likely to depend on immune cells, demonstrating that host-protective function acts directly in infected cell populations. Additionally, it's seen that the lethality in CASM-deficient mice derives from widespread uncontrolled viral entry and inflammatory signalling.

The function of the LIR motif remains obscure, however. As with canonical autophagy, it seems that if the LIR motif potentiates induction of CASM it will actively be limiting its own replication. It is possible that this is an evolutionary adaptation towards avirulence, and this hypothesis is consistent with the reports of combined increased pathogenicity and altered patterns of LC3 lipidation in highly pathogenic avian influenza strains. However, CASM seems to reduce influenza replication primarily via slowing virion-endosome fusion as described in (Yingxue Wang et al. 2021), and thus possibly acts in its capacity as host defence prior to any interaction of the M2 cytoplasmic tail with LC3. It is possible, then, that the potentiation of CASM by the LIR motif functions as a mechanism of superinfection exclusion (SIE), and thus is of primary epidemiological importance as SIE is a fundamental mechanism in limiting the rate of generation of reassortant strains (Sims et al. 2022). This is consistent with reports of higher multiplicities of infection (and thus greater CASM induction) generating stronger SIE effects (Martin et al. 2020). It is equally possible that the effectors of the CASM response in influenza infection are functionally saturated even in the absence of the LIR motif, and the increased

lipidation doesn't modulate the host response in a meaningful manner. Pointing towards this possibility is the effect of CASM on the viral lifecycle prior to viral entry and thus prior to substantial LC3 lipidation. Instead, the LIR motif may simply be a way for M2 to hijack on host trafficking systems, something that may be particularly important considering influenza's relatively unique independence from the host cell's ESCRT machinery (Rossman et al. 2010). Finally, it may present a way of interpreting cellular signals to allow influenza to adapt to its cellular environment, something that is rarely described in viruses but could represent an interesting mechanistic manifestation of a virulence-transmission trade-off on a per-cell level (Dolnik et al. 2015).

1.2.4 Conclusions

Measuring LC3 lipidation is not sufficient to differentiate canonical autophagy from CASM, and this error has led to a confused literature and misaddressed interpretations of experiments - in some cases with attendant clinical recommendations (Yan et al. 2013). Additionally, basic measures of LC3 relocalisation in microscopy images (such as spot count) are similarly non-specific to differing forms of autophagy. In order to more closely understand the relationship between M2 and CASM, it is necessary to develop more sophisticated analytical tools which are capable of identifying phenotypically distinct populations of cell images.

1.3 Observational and Quantitative Microscopy

1.3.1 Introduction

In fluorescence microscopy, it is possible to make a distinction between observational and quantitative biology (Wait, Reiche, and Chew 2020). Microscopes are fundamental tools for exploratory science, and descriptions from images entirely absent the use of statistical or quantitative methods are critical methods for hypothesis generation. Observational usage of

microscopes can be contrasted with the use of microscopes as quantitative tools, where a quantitative hypothesis is generated and measurements derived from images are used to test such a hypothesis. This is often seen in drug or genetic screening, where high-throughput microscopes take images of cells treated with a number of compounds so large so as to be infeasible to score by eye.

Both canonical autophagy and CASM are highly heterogeneous processes under most inducing conditions and are difficult to robustly differentiate using traditional fluorescence microscopy. This presents a difficulty to performing observational science as finding representative images amongst the phenotypic noise is difficult and highly prone to bias. However, as seen above, generating and testing valid quantitative hypotheses is ill-posed without robust knowledge of the phenotype we wish to measure. What is needed, therefore, is the application of quantitative methods to allow for observational analysis under significant biological variation and measurement noise. A section of the literature is concerned with pursuing such hypothesis-free data exploration in the context of high-throughput microscopy datasets (Chessel and Carazo Salas 2019; Piccinini et al. 2017). However, they often focus on datasets where there are significant and obvious phenotypic differences, and the advantage of quantitative data exploration techniques derives from their capacity to form observational inferences at the scale of thousands of different conditions. In this thesis, by way of contrast, the advantage of data exploration derives from ordering and controlling for noise in cellular populations, over a smaller number of selected, hypothesis-relevant conditions. This makes the analysis pursued in this thesis somewhat novel when situated in the wider literature, as we are using highly quantitative methods to make inferences that are as qualitative as those that would be seen in an observational microscopy study. Additionally, combined observational-quantitative approaches can be (and often are) as problematic as observational microscopy experiments with regards to

experimenter/analyst bias (Yanai and Lercher 2020), a consideration taken into account during the implementation and execution of this analysis.

1.3.2 Imaging flow cytometry

Imaging flow cytometry is a high-throughput imaging technique that captures images of cells in suspension as they flow through a microfluidics system that isolates and orders them within a single stream (Basiji et al. 2007). The primary advantage of this system is that it generates images of single cells, unlike other methods of high-throughput imaging which require segmentation prior to the generation of per-cell measurements. More subtle is the advantage that each cell image is acquired by itself, preventing complex batch-level correlations between cells segmented from the same image in plate-based automated imaging platforms. However, in comparison to confocal microscopy, imaging flow cytometry possesses significantly lower resolution, and requires cells to be in suspension prior to running, thereby resulting in a discrepancy between the physiological and acquisition environment for adherent cells. Taken together, the system is ideal for generating hypotheses about populations of single cells where phenotypes are not reliant on sub-micron detail, or adherent cell morphology, to be apparent.

1.3.3 Classical and Deep Feature Extraction

To pursue hypothesis-free image analysis, we require quantitative descriptors of images in our dataset. The evident difficulty is in choosing and designing our descriptors to be absent of analyst-introduced bias while still allowing for the focussed capturing of biologically relevant data. The relative advantages of traditional feature descriptor methods as compared to more novel deep learning techniques are difficult to determine *a priori*, especially as the literature concerned with generating feature descriptors is often situated in the context of natural images, which possess radically different statistical and pixel-level properties as compared to microscopy images.

Classically, generating representations requires a significant degree of domain expertise, especially in bioimage analysis. This is because the classical pipeline involves designing algorithms to extract features that will be of known biological interest (Handfield et al. 2013). Evidently, this requires not only knowledge of those features which are or are likely to be of interest but also of the image processing techniques that will best isolate those - often vaguely defined - features of interest. This results in a significant degree of variance between image analysts and thus introduces bias as to which features are emphasised and the quantitative form they take. Additionally, it limits the analysis to that of a quantitative assay and limits the possible extent of exploratory data analysis and thus the likelihood of finding novel phenotypes in a particular imaging dataset. The tradeoff is that because the features are hand-designed the resultant analysis is highly transparent and individual features are both relevant and easily interpretable.

Another approach to feature extraction in this context is to abandon handcrafted features and to use feature descriptors, developed early in computer vision research to classify datasets of natural images (Haralick, Shanmugam, and Dinstein 1973; Lowe 2004). Predictably, there is less direct analyst input and thus less introduced bias, but the sensitivity of the resulting analysis is likely to be low for biological features.

Methods to extract quantitative features using deep learning fall under the category of representation learning (Bengio, Courville, and Vincent 2013). The isolation of representation learning as an area of research is something of an arbitrary distinction, as all deep learning methods function by generating representations of data which are informative for a given task , most often classification. Representation learning, as a field, differs in the scope of its aims - where the representations of data are not only useful for the defined training task but also

possess a variety of desirable statistical properties. These properties often include interpretability and statistical independence of the individual quantities found in these representations (Yixin Wang and Jordan 2021), and allow them to function across a wide range of downstream tasks.

1.4 Aims

A traditional image analysis workflow for phenotypic profiling is predicated on strong prior knowledge of an imaging phenotype. After acquiring a large number of images, an image analyst would design image features with the intention that they accurately quantitate our known phenotype of interest, and subsequently judge the quality of these features by assessing how well they agree with our prior biological understanding of how our phenotype varies across a set of conditions.

A primary aim of this thesis is to extend this analysis to contexts where our knowledge of the phenotype is less strong and therefore is less suitable as an epistemic basis for developing image analysis pipelines. We anticipate this to occur in contexts where the biology itself is novel and thus knowledge of defining aspects of the phenotype is less secure. Additionally, if the phenotypic response is stochastic, or if the imaging technology introduces strong measurement noise, the increased heterogeneity within the acquired populations will make it difficult to assess the quality of any individual engineered feature. This analysis is designed to answer two questions of our data under these adverse inductive conditions. Firstly, to determine when two populations of cells are biologically distinguishable, and secondly to determine what these distinguishing phenotypes consist of. Influenza-induced LC3 lipidation is an ideal context for this analysis as the literature has clearly demonstrated a need for an approach to phenotypic profiling that can allow for the identification and separation of different LC3 patterning phenotypes.

2 - Results & Discussion

2.1 Unsupervised representation learning and high-throughput image analysis

2.1.1 Introduction

The training tasks used for training neural networks vary widely, and the choice of task will radically alter the representations and thus information derived from an imaging dataset (Hermann and Lampinen 2020). A commonly used task is classification, where over the course of training representations are extracted which will best allow for the classification of input images. At the end of the training, the output classifications are discarded and the representations are retained for further analysis - as seen in (McQuin et al. 2018). However, if populations of images with a given label do not all possess identifying image features (as is often the case in biological image datasets, either due to measurement noise or heterogeneity inherent to the biological process under investigation) then there is an incentive for the network to extract irrelevant features that spuriously allow for the separation of labelled populations (Geirhos et al. 2020). In this way, the training task introduces a structure to the representation space that is non-trivially entangled with the actual underlying population structure in the dataset.

To circumvent this problem, a set of techniques that do not use labelling information can be implemented, typically referred to as unsupervised learning approaches (Radford, Metz, and Chintala 2015; Locatello et al. Jun 2019; Le-Khac, Healy, and Smeaton 2020). As these approaches are unable to rely on label data as a source of information, the tasks by which their performance (and thus quality of their representations) is measured are often more complex. I

implemented two deep learning models, an Information Maximizing Generative Adversarial Network (InfoGAN) (Chen et al. 2016), and the Bootstrap Your Own Latent model (BYOL) (Grill et al. 2020), as representatives of two broader paradigms of unsupervised representation learning - generative approaches and self-supervised approaches respectively (Liu et al. 2021). The following analysis was intended to determine how two well-understood, unbiased, deep learning methods would function in the context of bioimage datasets, and in particular how well they could be used as a tool for scientific inference.

2.1.2 Representation learning with the InfoGAN

The InfoGAN relies on the training of a generative adversarial network at its core. Generative adversarial networks were - at the time of investigation - the gold standard for synthesising photorealistic images (Karras et al. 2019) (although they have since been superseded by methods such as stable diffusion (Rombach et al. 2021)). It was reasoned that their capacity to produce detailed images (as compared to other generative methods such as the variational autoencoders, which are notoriously blurred (Dosovitskiy and Brox 2016)) would allow them to better capture the relatively subtle texture variations that can define phenotypic differences in bioimage datasets. This has led to their previous use in fluorescence microscopy contexts (Goldsborough et al. 2017). The overall model consists of two individual networks in competition (hence the adversarial nature of the training), where one network acts as a generator and its companion as a discriminator. Training proceeds by allowing the generator to output a set of images, and allowing the discriminator to classify a mixed set of generated and real images into those which come from the true data distribution. The weights of the discriminator are subsequently updated so that it can better identify fabricated images, and the weights of the generator are updated so that it can more easily fool the discriminator network. As the networks are trained, the discriminator becomes more discerning, thereby providing a more stringent

training signal to the generator, and thus over time increasing generated image quality (K. Wang et al. 2017).

As neural networks are deterministic with regards to their input, we need to provide a form of random noise to the generator so that it can generate images of sufficient diversity. In the traditional GAN setup, this noise is of limited importance (although the exact form of the distribution can alter stability of the training process). However, in the InfoGAN an additional constraint is placed on the network, where the discriminator is also trained to infer the exact noise values fed into the generator. As the generator only has access to the images, it must infer these values solely from the image data, and thereby forcing the network as a whole to relate the most important features to these input noise values. This constraint is intended to force representations to consist of only those features which are most important for image generation, and thus exclude features which typically contribute to measurement noise - such as the position and orientation of the cell in the image (Kingma and Welling 2019).

2.1.3 Representation learning with BYOL

Self-supervised image representation learning (of which BYOL is an example) begins by first defining a set of programmatic image manipulations which do not fundamentally alter the semantic content of an image. These content-preserving manipulations are known as augmentations, and typically consist of spatial transformations such as rotation or flipping, or the application of basic image filters such as a Gaussian blur (Shorten and Khoshgoftaar 2019). Training proceeds by presenting two sets of images to the neural network, each of which has had a random subset of the defined augmentations applied. The network correspondingly outputs two sets of representations, and is trained such that the two sets are as similar as possible. Therefore, the network is trained to generate representations which are implicitly insensitive to the augmentations we define as content-preserving (von Kügelgen et al. 2021).

Such networks suffer from a problem known as “representational collapse”, whereby the training criteria of having identical representations under transformation is trivially met by the network outputting identical representations for every image. Typically, this is solved through the use of negative, or contrastive samples (Jing et al. 2021). This is where we additionally stipulate that a given representation must not only be insensitive to augmentation but also must be separable from a representation derived from a different image - the negative sample. This introduces a choice as to the method of selection of negative samples, a choice which can alter the form of extracted representations in ways which are dependent on the relative distribution of features in a given dataset (M. Wu et al. 2020). BYOL implements a different approach to this problem, in which no negative samples are used and representational collapse is prevented utilising two networks for each of the two augmented image sets. One network is set to be an exponential moving average of the other, and this results in relative stability during training - although the mathematical reasoning behind exactly why this approach prevents representational collapse is contested (Richemond et al. 2020; Fetterman and Albrecht 2020; Tian, Chen, and Ganguli Jul 2021). This lack of need for a negative sampling strategy was reasoned to be an advantage for the aims of this research as I wished to make as few assumptions about the distributions of features across the acquired dataset as possible.

2.1.4 Results

To generate the dataset, four EGFP-LC3 expressing HCT116 cell lines were transfected with four different M2 constructs under Tet-On inducible control - a wild-type construct from the Udorn strain of the virus, and three mutations of interest. The fluorescently-tagged LC3 allows us to monitor its relocalisation under different autophagic processes, and the inducible control allows us to induce M2 expression, and through M2 induce LC3 relocalisation, using doxycycline treatment.

The first mutation consists of a mutation in the LIR motif, F91S, known to significantly attenuate the extent of the M2-induced LC3 lipidation. Two additional mutants are less well characterised but are of interest - the L4Q mutation which occurs in the N-terminal ectodomain and may alter the localisation of the M2 protein (Wise et al. 2012), and the R45H mutation which is found in reversion mutants of strains with highly sensitive haemagglutinin fusion machinery, implying either an altered ion channel or an altered channel gating mechanism to more strongly raise pH in the TGN.

Each of these cell lines was treated with doxycycline at 5 µg/mL for 16 hours to induce expression of the M2 protein, alongside a mock condition to capture the cell lines at a baseline with no CASM response. After trypsinisation, cells were run unfixed on an imaging flow cytometer to acquire fluorescence images, thus only capturing the EGFP-LC3 fluorescence channel. After neural network training, features were extracted with the InfoGAN by passing images to the discriminator and extracting activations from its final linear layer, as this seems to drive improved performance. After training with the BYOL, images were passed to the ResNet backbone and features were extracted from the final maxpool layer.

Upon immediate inspection, features were difficult to parse - seemingly admitting quite a significant degree of variation between control populations. As neural networks were intended to capture a large proportion of the variation in image datasets, I adapted a method of feature normalisation from the literature to more closely align the baseline phenotypes across cell lines which is described in figure 3. Briefly, a robust measure of standard deviation for each feature is calculated across firstly all control populations, and subsequently across the total dataset. If a feature varied minimally within the control populations but substantially across the dataset it was reasoned that this would indicate the feature would likely relate in some way to the expression

of M2 and thus be of biological interest. The ratio of control to total variation, therefore, was used as a method of ranking and filtering the extracted features.

The features from both deep learning approaches were analysed and ranked as above, and low ranking features were progressively ablated prior to undergoing a principal components analysis to determine how amenable the respective extracted representations were to the removal of nuisance variation (figure 1A). This progressive ablation performed for the InfoGAN features shows a limited effect on the degree of variation seen across control populations, but as performed for the BYOL-extracted features the variation in control populations is seen to be strongly controlled. Additionally, when the 1000 highest ranked features are kept and the differing control conditions along the first principal component axis are plotted (figure 1B), it can be seen that the BYOL features show the greatest consistency across control conditions, whereas the InfoGAN features display significant heterogeneity across control conditions.

To determine whether the resultant principal component axis corresponded to biological variation, an interpolation along this axis was generated, and images close to those interpolated values were plotted (figure 2B). Both axes corresponded at least loosely to increasing relocalisation of EGFP-LC3, with the particular perinuclear relocalisation seeming to be in agreement with the known phenotype of CASM. However, the InfoGAN representations seemed to be closely following the mean fluorescence of the images, with the BYOL representations admitting more variation in similarly basic features. To further explore these visual distinctions, 64 images from the first and last deciles of the InfoGAN and BYOL first principal components were plotted (figure 2A). It can be further seen that the InfoGAN representations are very strongly dependent on image fluorescence, whereas the BYOL features demonstrate a strong distinction primarily in patterning.

To quantitate the dependence of the InfoGAN features on photometric features, I performed a basic second-order polynomial regression of the first principal component against fluorescence intensity (figure 2C). It was seen that the InfoGAN principal component much more closely follows intensity values, whereas the BYOL features seem to admit a much greater degree of variation of intensity values when extracting phenotypic information.

Finally, to determine whether these principal component measures could accurately determine the phenotypic differences between mutants, we constructed a measure of CASM induction based on the Jensen-Shannon divergence (JSD) (figure 2D). This is a well-known measure of the difference between two distributions (Menéndez et al. 1997), and it was reasoned that the difference between control and induced populations would be a better measure of phenotype induction than a simple mean of the induced populations. This is because it would in theory take into account any differences found at baseline in the different cell lines. Unexpectedly, the InfoGAN first principal component axis shows closer agreement with our previous knowledge of the phenotype, where the F91S significantly ablates the induction of CASM. This is in contrast with the BYOL embeddings, which seem to imply no significant difference between the M2 WT and F91S mutants.

2.1.5 Discussion & limitations

The primary aim of these results were to investigate the performance of neural networks in a targeted phenotypic environment. The reasoning behind the use of neural networks was that it would remove a degree of bias from attempts at visual phenotyping in relatively novel contexts, such as underexplored mutations in a viral protein. By attempting to make best use of neural networks in this context, however, some aspects of this reasoning have needed reexamination.

Firstly, as the degree of information contained with a neural network representation greatly exceeds the real degree of phenotypic variation in our particular small-scale datasets, data pruning is required to remove the nuisance sources of variation the networks will inevitably model. The difficulty is that smaller scale networks with a more limited encoding capacity do not model complex phenotypic details particularly well. Choosing or designing methods, as I have done in this context, to isolate signal from noise in the resultant representations will consequently generate similar levels of analyst-derived bias as would hand-designing image features. Additionally, the normalisation analysis (and those analyses which it is based on) rely on there being minimal actual phenotypic variation across our control populations - an assumption that could be easily violated due to clonal artefacts altering the baseline degree of M2 expression, and is often seen in other high-throughput contexts.

Critically, the multiple complex choices involved in not only designing an architecture for a neural network, but performing the necessary data engineering to allow the input images to be suitable for training, and the choices of training parameters, will all substantially impact the form and content of representations extracted by using neural networks - as demonstrated by the significant differences in the distributional properties of the InfoGAN and BYOL representations. These choices that a neural network practitioner must make alter the inductive bias of such networks, choices which are as substantial and less well understood than those made to algorithmically extract known features from images.

An example of the difficulty that this introduces to attempting to derive scientific knowledge from these features can be seen with the BYOL representations and their independence from fluorescence intensity. One of the augmentations we used as part of the self-supervised training was that of altering the fluorescence of the images. This was intended to drive an inductive bias in the network whereby it would learn that parts of the semantic content of an image were

independent of fluorescence intensity. It is unclear whether this independence from intensity improves representations because intensity is a source of nuisance variation, or whether it is a source of actual biological variation that can be ablated to ensure that control populations look distributionally similar. The frequency and degree with which this augmentation is applied during training will alter the degree of control over intensity in the representations. In this way, the decision to try and control for fluorescence intensity is not only made implicit but also more difficult to justify, interpret and, most importantly, critique.

Collectively, these considerations prompted some questioning as to the actual utility of high-throughput imaging in this particular context of targeted phenotypic discovery. The original aim of generating a set of representations which immediately separate populations with minimal processing or encoded prior phenotypic knowledge was likely to be ill-founded. It was reasoned that there is enough spurious information contained within any individual condition to allow it to be easily classified as separate from another condition in the absence of actual phenotypic difference. Therefore, considering this sensitivity of feature extraction to nuisance variation, it would be possible to more concretely reason from extracted image features if two populations - known to be phenotypically identical - were indistinguishable under active attempts at separation in a given feature space. Validating and performing scientific inference in representational spaces using this indistinguishability criterion forms the basis of the following analyses.

2.2 Phenotypic identification with classical image features

2.2.1 Differentiating canonical and non-canonical autophagy

To investigate this alternative approach to phenotypic identification in bioimage datasets, we aimed to try to identify and retrieve the known phenotypic differences between canonical and non-canonical autophagy.

To induce canonical autophagy we treated cells with both Torin-1, an mTOR inhibitor, and bafilomycin A, a vATPase inhibitor that is thought to inhibit the transduction of CASM-inducing signals, thus generating a solely canonical phenotype (Florey et al. 2015). To induce CASM, we treated HCT116 cells with monensin, an ionophore which disrupts ion gradients across a wide range of compartments and, in a similar way to M2, drives the vATPase-mediated recruitment of ATG16L1 that underlies the CASM response. However, due to the breadth of the targets of monensin in the cell, the drug is toxic and can induce several changes in cell morphology even in the absence of CASM. To deconvolute the CASM specific responses in phenotype from those due to the broadly deleterious effects of monensin, these treatments were applied in two genetic backgrounds - a CRISPR KO of ATG16L1, reconstituted with either wildtype ATG16L1 or a construct with a point mutation in the WD40 domain of the protein, K490A, which has been shown to strongly (albeit not completely) ablate the capacity of the cell to mount a CASM response while having no impact on canonical autophagy (Fletcher et al. 2018). In the resultant population of cells, it was reasoned that only those cellular responses unique to the monensin-treated, wildtype-reconstituted cell line would represent the true CASM imaging phenotype.

Three biological replicates of each condition were fixed, stained and run separately, with 10,000 cell images captured for each replicate. The cells were fixed and stained for LC3, and TGN46,

as a *trans*-Golgi marker which - at the resolution obtained - acts as a readout for the perinuclear area. After filtering for image quality, a set of 126 image features was extracted using the IDEAS software.

By eye, it was found that a colocalisation measure of LC3 with TGN46 was the best individual image feature for differentiating the WT Monensin population from other conditions (figure 4H). In agreement with this observation, a key individuating characteristic of monensin-induced CASM is relocalisation to the perinuclear area. However, it was seen that high colocalisation values could not uniquely identify the WT Monensin population, as high values of colocalisation could be found in all populations (figure 4I).

To find the phenotypically unique population of cell images in the WT Monensin condition, we used linear discriminant analysis (LDA), which analogously to principal component analysis (PCA) projects high dimensional data into a set of summary variables (figures 4C, 4D, 4E). LDA differs from PCA in that it uses class information to project the data into summary variables which best separate the given classes, as opposed to the variance-maximising projections of PCA (Hastie, Tibshirani, and Friedman 2009) pp. 106 - 119. We perform this analysis to explicitly attempt to distinguish between the WT and K490A populations under the mock and Torin/Baf treatments where they should be identical. We can then measure distributional distances between conditions in the first two dimensions of the space to determine whether they can be reasonably differentiated (figure 4B). The distance metric we use in this case is the Wasserstein distance, as it possesses better numerical stability properties (Arjovsky, Chintala, and Bottou 2017). It is seen that the distributional differences between repeats are on a similar scale to the difference between genotypes under mock and Torin/Baf treatments, suggesting that the representational space that we have extracted is appropriately insensitive to nuisance variation without the need for complex normalisation strategies.

Under monensin treatment the wildtype and K490A cell lines are highly separable (figure 4D), showing that the two cell lines are only distinguishable under conditions where CASM is induced. To further investigate the performance of this analysis, the data were projected onto the axis that best separated the WT Monensin condition from the other conditions (figures 4F, 4G). Evidently, this provided much better separation from the other cell images than the naive single-feature approach, with a proportion of the WT Monensin population being uniquely attributed to the CASM phenotype. Using this WT Monensin axis to generate a cutoff, we can compare images from the population that pass the cutoff to those that don't (figure 4A), as a form of control internal to the WT Monensin population. We can subsequently visually identify features which seem to correspond to the induction of CASM, namely the diffuse spherical localisation of LC3 to a single site in the cell, often perinuclear. This can be compared to the cells which show no discernible LC3 patterning (suggesting a lack of atg8-ylation to membranes) or a multilobed/punctate LC3 pattern that does not colocalise with the perinuclear marker, possibly indicating a canonical autophagic response.

2.2.2 Investigating influenza M2 mutants

In the knowledge that the representational space extracted using basic image features seemed to be reasonably robust to nuisance variation, we used a similar analysis to again investigate the phenotypes of the influenza M2 mutants, with an experimental design that allowed for the statistical analysis of distributional distances, where we generated three biological repeats for each condition.

We transduced three M2 constructs as before into a polyclonal rather than monoclonal EGFP-LC3 background - the wildtype construct, and two mutants, F91S and R45H. As before, three biological replicates of each condition were fixed, stained and run separately, with about

1500 cell images captured for each replicate. The cells were fixed and stained for M2, to allow for filtering for the M2 positive populations. Additionally, as there was seen to be a strong division of the M2 protein between the plasma membrane and perinuclear compartments, features were calculated over cytoplasmic and perimembrane masks in addition to the whole cell mask. After filtering for image quality as before, a set of 259 image features was extracted using the IDEAS software.

To determine whether the populations were phenotypically unique, we performed the LDA analysis as before (figure 5C) seeing that all different mutants were strongly distinguishable. To determine the relative contribution of the two fluorescence channels to this distinguishability, we again performed an LDA analysis but using only those features derived from either the EGFP-LC3 channel or the M2 channel (figure 5E). Direct quantification of distributional distances between conditions demonstrates that the EGFP-LC3 pattern is the primary distinguishing factor for the F91S mutant, whereas both M2 and EGFP-LC3 patterning contribute to the separation of R45H from the other conditions (figure 5F). To determine the exact phenotype contributing to the discrimination of R45H from other conditions, we can determine, as before, the axis of greatest discrimination for the R45H condition (figure 5B). Visualising samples from the extremes of the R45H population along this axis, it can be seen that there is lesser M2 intensity and seemingly greater EGFP-LC3 relocalisation (figure 5A), and this difference can be quantified (figure 5D). This suggests that within a heterogeneous population of EGFP-LC3 relocalisation induction, those cells which induce strong relocalisation at lower levels of M2 expression are likely to be those cells expressing the R45H mutant. This is consistent with the proposed greater degree of ion conductivity with this mutant.

As the F91S mutant populations display little overlap with the other two conditions (figure 5C), isolating the unique portions of this population is unnecessary. Instead, we can easily visually

identify that the M2 population seems to have a greatly reduced relocalisation of EGFP-LC3, and seemingly very little M2 localised intracellularly. To quantify these patterns, we use a measure of bright detail intensity (which filters for high frequency patterns before measuring intensity of an image, figure 6C) and the modulation of M2 in a cytoplasmic mask (which effectively calculates the maximum pixel range, figure 6D). With this quantification, it is seen the distribution of these patterns is greatly altered across mutants (figures 6A, 6B, 6E). This seems to indicate that in addition to driving differing degrees of CASM induction, the different mutants - possibly as a result - also undergo altered trafficking through the cell.

To determine whether this alteration in M2 localisation was an artifact induced by high non-endogenous expression of EGFP-LC3, for example by saturating the LIR motif binding site in a non-physiological manner, the difference in M2 modulation was measured for differing quintiles of EGFP-LC3 expression (figure 6F). This seemed to demonstrate that the difference between mutants in M2 localisation was preserved across all levels of EGFP-LC3 expression. However, there was seen to be an increase in the degree of M2 internal localisation for the wildtype M2 construct as GFP-LC3 expression increased. Due to the permeabilisation process, and despite strong formaldehyde fixation, some proportion of unlipidated LC3 will be removed from the cell - causing a potential correlation between GFP-LC3 intensity and the degree of its lipidation (Eng et al. 2010). Thus the increase of M2 internal localisation with GFP-LC3 intensity could possibly reflect a dose-dependent effect of M2 expression on the degree of LC3 lipidation.

2.2.3 Delineating effects of ion channel activity and LC3-binding on M2 and LC3 patterning

There are two possible explanations for the difference in M2 localisation between the wildtype and F91S cell lines. The first is that the M2 protein is reliant on CASM for the modulation of its trafficking, and the reduction in CASM seen with the F91S mutant is directly responsible for its

altered localisation. The second is that the LC3-binding capacity of M2 is directly causally linked to its trafficking, and the lack of a functional LIR motif is responsible for both the attenuated CASM response and the altered trafficking of the F91S mutant. To investigate this, we used amantadine as a tool to block M2 ion channel activity, thereby dissecting the two contributors to the combined trafficking and CASM response - the LC3-binding capacity and ion gradient disruption of the protein.

Four conditions were investigated: wildtype and F91S with either only doxycycline treatment or doxycycline and amantadine. Three biological replicates of each condition were fixed, stained and run separately, with approximately 3500 cell images captured for each replicate. The cells were stained for M2 as before, to allow for filtering for the M2 positive populations, and features were calculated over cytoplasmic and perimembrane masks in addition to the whole cell mask. After filtering for image quality, a set of 380 image features was extracted using the IDEAS software.

An LDA analysis was performed, and all conditions were seen to be separable, with the strongest phenotype seen to be WT untreated with amantadine (figure 7A). As before, the WT condition demonstrated both the greatest degree of EGFP-LC3 relocalisation and M2 internal localisation (figure 7B, 7C). To delineate effects on EGFP-LC3 relocalisation and M2 localisation we performed a further set of LDA analyses, firstly where solely those features extracted from the LC3 channel were used as input features, and secondly where only those features extracted from the M2 channel were used (figure 7D, 7F). These two analyses demonstrated differing patterns of distinguishability. It was seen that for the M2-derived features, there was no discernible difference between F91S with and without a functional ion channel, but a statistically significant difference between WT and F91S when both were treated with amantadine (figure 7E). Conversely, the same statistical analysis performed with the LC3-derived features

demonstrated the opposite pattern of phenotypic separation (figure 7G). Taken together, this is consistent with the primary dependence of CASM on the presence of a functional ion channel as the initiating stimulus, and that the lack of a LIR motif simply prevents potentiation - as F91S mutants still demonstrate some degree of CASM induction. Additionally, in the absence of detectable CASM under amantadine treatment, it is seen that there must be some biological interaction of M2 with LC3 as there is a discernible trafficking difference in the WT Dox/Amt population. It is possible that this results from reduced amount of breakthrough CASM induction that is undetectable with our assay, or alternatively that M2 is capable of a trafficking-relevant interaction with LC3 in the absence of a major lipidation stimulus.

2.3 Discussion & Future work

I hope to have demonstrated that this form of phenotypic analysis allows us to generate novel biological insights, not through the synthesis of highly diverse phenotypic data as might be seen in a high throughput drug screen, but rather through highly targeted understanding of a smaller set of phenotypes. In this targeted experimental regime the high data volume allows us to make strong inferences about perceptual difference which would otherwise be difficult due to inherent biological heterogeneity and measurement noise. It additionally allows for the statistical control of potential overexpression artifacts, and overall allows us to some extent mitigate the limitations of the experimental tools we use in the imaging of autophagic phenotypes.

The above analysis has provided several biological insights. The first is that M2 is dependent on its interaction with LC3 for its trafficking through the cell. This had never been found in previous descriptions of LC3 relocalisation under M2 expression (Beale et al. 2014), likely because attempting to infer this insight from lower throughput data would have been difficult. Firstly, as the phenotypic overlap seen with the simple measure of M2 internal localisation (figure 6B) would have lessened the perceptual difference needed for concrete conclusions from confocal

microscopy data. Additionally, controlling for EGFP-LC3 overexpression artefacts via the quintile analysis would have equally been non-viable without several thousand cell measurements. A further insight is that M2 has a measurable interaction with LC3 in the absence of substantial CASM induction. Similarly, quantifying the perceptual differences which underpin the dissection of the contribution of the LIR motif and ion channel activity to M2 localisation would have been infeasible without the utilisation of measures of distributional difference to quantify the perceptual dissimilarity specific to each channel. Attempting to prove or disprove a functional biological interaction of M2 with LC3 in the absence of CASM would otherwise be very difficult to prove with biochemical techniques.

Finally, the use of LDA to identify populations unique to a given class allowed us to firstly identify a LC3 patterning signature unique to CASM despite significant phenotypic inference from monensin, and secondly identify that the R45H M2 mutant was able to induce CASM at lower expression levels than a wildtype M2 protein. Notably in this latter case the patterning in both channels was similar to the wildtype, and the identification of this increased potency derived from an inference in the *relationship* between the two channels, suggesting that this form of analysis can allow for the discovery of subtle mechanistic differences in phenotypes.

Notably, in comparison to the deep learning techniques covered earlier, the features which were used did not undergo significant processing to control for nuisance variation, suggesting that the relatively simple extracted image features that were used did not extract enough noise to generate spuriously unique populations. This, alongside the significantly reduced implementation time and compute usage, represents a significant advantage of feature extraction over representation learning in this particular context.

There are two strong biological limitations in these experiments which restrict the breadth of conclusions we can draw from these results. The first is the use of an M2 protein expression system rather than directly studying influenza A infection. This may not recapitulate the trafficking pathway of the M2 protein during infection, due to the expression of other viral proteins that can interact with M2 such as M1. This is compounded by the fact we have only studied this phenotype in the HCT116 cells, and may rely on some idiosyncratic aspect of this cell line to manifest. The strength of the effect and its response under amantadine treatment argue against this, however.

If this patterning phenotype is insensitive to cell line background, it may provide a provocative indication of the function of the LIR motif. CASM is protective for the host in a mouse model (Yingxue Wang et al. 2021). However, the fact that there is a conserved LIR motif within the M2 protein suggests that there is some paradoxical benefit to the virus to not only induce - with the action of the ion channel - but seemingly amplify this host protective response. If the presence of LC3 acts as a form of trafficking regulation for M2, this could represent a system by which influenza alters its lifecycle in response to host signalling. There is evidence to suggest that M2 is required for budding of virions from the plasma membrane (Rossman et al. 2010; Rossman and Lamb 2011) and from this it is possible to speculate that the concentration of M2 at the plasma membrane would be regulated to ensure proper co-ordination of the budding machinery. Notably, influenza virions in their natural state are highly pleiomorphic (Stevenson and Biddle 1966), and demonstrate a wide distribution of shape from effectively spherical to filamentous. Mathematical modelling of this pleiomorphism has indicated that it may serve an important function in adapting to host humoral immune responses (T. Li et al. 2021), with different shapes being optimal at different levels of antibody prevalence. It is therefore possible that different levels of CASM alter virion shape by differentially retaining M2 within the cytoplasm, and that this system represents a method by which influenza alters its reproductive strategy in response

to cellular stress. This theory is consistent with evidence demonstrating that F91S M2 mutant viruses display a strongly limited ability to generate filamentous virions (Beale et al. 2014).

Further computational work should investigate whether altered architectures and training paradigms (for example, tailoring the augmentations used for contrastive learning to bioimaging contexts) can improve not only the information content of the resultant representations, but also their amenability to direct interpretation by simple, linear methods such as LDA. Critically, metrics that go beyond downstream classification accuracy should be developed to capture non-standard dimensions of representational performance. In contexts where reproducibility and visual interpretability are key, such metrics would help improve the transparency of architectural and analytical decisions made by deep learning practitioners. Closer inspection of the ways in which such decisions alter the form and content of scientific inferences based on these deep learning models will, in turn, allow for much needed epistemic trust in their outputs.

3 - Materials & Methods

3.1 Transfection & Generation of M2 Mutant Cell Lines

M2 and associated mutant constructs were inserted into a pInducer20 plasmid - a gift from Stephen Elledge (Addgene plasmid # 44012 ; http://n2t.net/addgene:44012 ; RRID:Addgene_44012). Second generation lentiviral vectors were generated in HEK293T cells, under polyethylenimine transfection. Polyclonal EGFP-LC3 HCT116s were transduced with lentivirus and passaged under G418 selection at 400 μg/mL for 5 days, until untransduced cells were no longer detectable under brightfield microscopy.

3.2 Imagestream Cell Preparation

HCT116s are seeded at a density of 20000 cells per well in 12 well plates and grown for two days at 37ºC with 5% CO_2 prior to treatment. Wells are treated as indicated in the text. After treatment, wells are washed once with 4ºC PBS, and cells are dissociated using 800µL of Accumax (ThermoFisher Scientific). The dissociation reaction is quenched with 200µL of 4ºC 5% FCS in PBS, volumes are transferred to 1.5mL eppendorf tubes and cells are spun for 5 minutes in a 4ºC microfuge at 200g. The dissociation reagent is aspirated and cells are further washed in 4ºC PBS and spun. Cells are resuspended and fixed in 4% formaldehyde at 4ºC for 40 minutes. Cells are washed and resuspended in a 0.1% saponin, 5% FCS in PBS permeabilisation/staining buffer. Cells are stained with primary antibodies at 4ºC overnight, and secondary fluorescent antibodies for one hour at room temperature. Cells are resuspended in PBS prior to running on the Amnis ImageStream X Mk II (Luminex).

3.2.1 Imagestream Cell Preparation - Antibodies

Mouse anti-LC3 (Cosmo Bio - Clone LC3.No.6, 1:100) was used to stain for LC3 patterning. Mouse anti-M2 (Abcam - ab5416, 1:200) was used to stain for M2 patterning. Rabbit anti-TGN46 (Abcam - ab50595, 1:1000) was used to stain for M2 patterning. Llama anti-mouse CF488A (Biotium, 1:1000) was used to label LC3, Goat anti-Rabbit AF647 (ThermoFisher Scientific, 1:1000) was used to label TGN46, and finally Goat anti-Mouse AF647 (ThermoFisher Scientific, 1:1000) was used to label M2.

3.3 Imagestream Data Acquisition and Analysis

Cells were run for approximately 15 minutes per sample. Acquired populations were gated to exclude speedbeads, cell aggregates, and cells out of the plane of focus. Laser power for each channel was set to prevent detector saturation.

After acquisition, cell images where cells clip the image border are removed using a metric of circularity on the brightfield image mask. The IDEAS software (Luminex) is used to extract a suite of image texture features and mask shape features. Additionally, cytoplasmic masks are generated via use of an intensity weighted pixel-wise erosion function on a mask calculated from a brightfield image. The perimembrane mask is generated by subtracting the cytoplasmic mask from the overall brightfield mask. Features are subsequently exported to text files for further analysis.

3.3.1 Data analysis

Full code to perform the classically feature extraction analyses can be found at https://github.com/ooakley/IS_Analysis - the exported feature values exceed the maximum repository size and are available upon request. Jupyter notebooks were used as an interactive coding environment, and the primary packages used were scikit-learn, seaborn, and scipy - full package requirements can be found in the conda environment file. PRISM (GraphPad) was used to plot summary features and perform statistical analyses.

3.4 Neural Network Methods

3.4.1 Neural Network Data Engineering

Raw images were exported directly from the Amnis compensated image files. Data preparation for fluorescence images was performed as previously described (Ísaksson 2018), briefly by

fitting a three component Gaussian distribution to the fluorescence pixel values across the whole dataset, and subsequently performing z-score normalisation using the mean and standard deviation of the component with the largest mean. Full code and computational details can be found at https://github.com/ooakley/ifc-data-engineering.

3.4.2 Neural Network Architecture and Training Setup

Repositories for the BYOL and InfoGAN architectures can be found at https://github.com/ooakley/byol-embeddings and https://github.com/ooakley/ReusableInfoGAN, respectively. The BYOL implementation is heavily indebted to an open source implementation developed by (P. Wang 2020).

The InfoGAN was trained for 20 epochs with a batch size of 100 (or approximately 73,000 gradient updates). 64 uniform dimensions, 2 categorical dimensions and 5 Gaussian noise dimensions as input to the generator network. Gradient regularisation as described in (Roth et al. 2017) was implemented, but due to the low-dimensional manifold of bioimages as compared to natural images, a form of instance normalisation known as spectral normalisation was needed to stabilise gradients (Miyato et al. 2018). The BYOL was trained for 200 epochs with augmentations as described in the given repository.

4 - Bibliography

Alvarado-Facundo, Esmeralda, Yamei Gao, Rosa María Ribas-Aparicio, Alicia Jiménez-Alberto, Carol D. Weiss, and Wei Wang. 2015. “Influenza Virus M2 Protein Ion Channel Activity Helps to Maintain Pandemic 2009 H1N1 Virus Hemagglutinin Fusion Competence during Transport to the Cell Surface.” *Journal of Virology* 89 (4): 1975–85.

Appleyard, G. 1977. “Amantadine-Resistance as a Genetic Marker for Influenza Viruses.” *The Journal of General Virology* 36 (2): 249–55.

Arjovsky, Martin, Soumith Chintala, and Léon Bottou. 2017. “WAsserstein Generative Adversarial Networks.” Edited by Doina Precup and Yee Whye Teh, Proceedings of Machine Learning Research, 70: 214–23.

Basiji, David A., William E. Ortyn, Luchuan Liang, Vidya Venkatachalam, and Philip Morrissey.

2007. “Cellular Image Analysis and Imaging by Flow Cytometry.” *Clinics in Laboratory Medicine* 27 (3): 653–70, viii.
Beale, Rupert, Helen Wise, Amanda Stuart, Benjamin J. Ravenhill, Paul Digard, and Felix Randow. 2014. “A LC3-Interacting Motif in the Influenza A Virus M2 Protein Is Required to Subvert Autophagy and Maintain Virion Stability.” *Cell Host & Microbe* 15 (2): 239–47.
Bengio, Yoshua, Aaron Courville, and Pascal Vincent. 2013. “Representation Learning: A Review and New Perspectives.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35 (8): 1798–1828.
Bergmann, M., A. Garcia-Sastre, E. Carnero, H. Pehamberger, K. Wolff, P. Palese, and T. Muster. 2000. “Influenza Virus NS1 Protein Counteracts PKR-Mediated Inhibition of Replication.” *Journal of Virology* 74 (13): 6203–6.
Calle, Claire de la, Pierre-Emmanuel Joubert, Helen K. W. Law, Milena Hasan, and Matthew L. Albert. 2011. “Simultaneous Assessment of Autophagy and Apoptosis Using Multispectral Imaging Cytometry.” *Autophagy* 7 (9): 1045–51.
Chen, Xi, Yan Duan, Rein Houthooft, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. “InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets.” *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/1606.03657.
Chessel, Anatole, and Rafael E. Carazo Salas. 2019. “From Observing to Predicting Single-Cell Structure and Function with High-Throughput/high-Content Microscopy.” *Essays in Biochemistry* 63 (2): 197–208.
Ciampor, F., P. M. Bayley, M. V. Nermut, E. M. Hirst, R. J. Sugrue, and A. J. Hay. 1992. “Evidence That the Amantadine-Induced, M2-Mediated Conversion of Influenza A Virus Hemagglutinin to the Low pH Conformation Occurs in an Acidic Trans Golgi Compartment.” *Virology* 188 (1): 14–24.
Comber, Joseph D., Tara M. Robinson, Nicholas A. Siciliano, Adam E. Snook, and Laurence C. Eisenlohr. 2011. “Functional Macroautophagy Induction by Influenza A Virus without a Contribution to Major Histocompatibility Complex Class II-Restricted Presentation.” *Journal of Virology* 85 (13): 6453–63.
Dolnik, Olga, Valentina A. Volchkova, Beatriz Escudero-Perez, Philip Lawrence, Hans-Dieter Klenk, and Viktor E. Volchkov. 2015. “Shedding of Ebola Virus Surface Glycoprotein Is a Mechanism of Self-Regulation of Cellular Cytotoxicity and Has a Direct Effect on Virus Infectivity.” *The Journal of Infectious Diseases* 212 Suppl 2 (October): S322–28.
Dosovitskiy, Alexey, and Thomas Brox. 2016. “Generating Images with Perceptual Similarity Metrics Based on Deep Networks.” *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/1602.02644.
Dubois, Julia, Olivier Terrier, and Manuel Rosa-Calatrava. 2014. “Influenza Viruses and mRNA Splicing: Doing More with Less.” *mBio* 5 (3): e00070–14.
Duve, C. de, B. C. Pressman, R. Gianetto, R. Wattiaux, and F. Appelmans. 1955. “Tissue Fractionation Studies. 6. Intracellular Distribution Patterns of Enzymes in Rat-Liver Tissue.” *Biochemical Journal* 60 (4): 604–17.
Eng, Kai Er, Marc D. Panas, Gunilla B. Karlsson Hedestam, and Gerald M. McInerney. 2010. “A Novel Quantitative Flow Cytometry-Based Assay for Autophagy.” *Autophagy* 6 (5): 634–41.
Feizi, Neda, Parvaneh Mehrbod, Bizhan Romani, Hoorieh Soleimanjahi, Taravat Bamdad, Amir Feizi, Ehsan Ollah Jazaeri, et al. 2017. “Autophagy Induction Regulates Influenza Virus Replication in a Time-Dependent Manner.” *Journal of Medical Microbiology* 66 (4): 536–41.
Fetterman, Abe, and Josh Albrecht. 2020. “Understanding Self-Supervised and Contrastive Learning with Bootstrap Your Own Latent (BYOL).” 2020. https://generallyintelligent.ai/blog/2020-08-24-understanding-self-supervised-contrastive-learning/.
Fletcher, Katherine, Rachel Ulferts, Elise Jacquin, Talitha Veith, Noor Gammoh, Julia M. Arasteh, Ulrike Mayer, et al. 2018. “The WD40 Domain of ATG16L1 Is Required for Its Non-Canonical Role in Lipidation of LC3 at Single Membranes.” *The EMBO Journal* 37 (4).

https://doi.org/10.15252/embj.201797840.
Florey, Oliver, Noor Gammoh, Sung Eun Kim, Xuejun Jiang, and Michael Overholtzer. 2015. “V-ATPase and Osmotic Imbalances Activate Endolysosomal LC3 Lipidation.” *Autophagy* 11 (1): 88–99.
Florey, Oliver, Sung Eun Kim, Cynthia P. Sandoval, Cole M. Haynes, and Michael Overholtzer. 2011. “Autophagy Machinery Mediates Macroendocytic Processing and Entotic Cell Death by Targeting Single Membranes.” *Nature Cell Biology* 13 (11): 1335–43.
Gannagé, Monique, Dorothee Dormann, Randy Albrecht, Jörn Dengjel, Tania Torossi, Patrick C. Rämer, Monica Lee, et al. 2009. “Matrix Protein 2 of Influenza A Virus Blocks Autophagosome Fusion with Lysosomes.” *Cell Host & Microbe* 6 (4): 367–80.
Geirhos, Robert, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. “Shortcut Learning in Deep Neural Networks.” *Nature Machine Intelligence* 2 (11): 665–73.
Goldsborough, Peter, Nick Pawlowski, Juan C. Caicedo, Shantanu Singh, and Anne E. Carpenter. 2017. “CytoGAN: Generative Modeling of Cell Images.” *bioRxiv*. https://doi.org/10.1101/227645.
Grill, Jean-Bastien, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, et al. 2020. “Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning.” *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/2006.07733.
Handfield, Louis-François, Yolanda T. Chong, Jibril Simmons, Brenda J. Andrews, and Alan M. Moses. 2013. “Unsupervised Clustering of Subcellular Protein Expression Patterns in High-Throughput Microscopy Images Reveals Protein Complexes and Functional Relationships between Proteins.” *PLoS Computational Biology* 9 (6): e1003085.
Haralick, Robert M., K. Shanmugam, and Its ’hak Dinstein. 1973. “Textural Features for Image Classification.” *IEEE Transactions on Systems, Man, and Cybernetics* SMC-3 (6): 610–21.
Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning*. Springer New York.
Hermann, Katherine L., and Andrew K. Lampinen. 2020. “What Shapes Feature Representations? Exploring Datasets, Architectures, and Training.” *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/2006.12433.
Holsinger, L. J., M. A. Shaughnessy, A. Micko, L. H. Pinto, and R. A. Lamb. 1995. “Analysis of the Posttranslational Modifications of the Influenza Virus M2 Protein.” *Journal of Virology* 69 (2): 1219–25.
Hooper, Kirsty M., Elise Jacquin, Taoyingnan Li, Jonathan M. Goodwin, John H. Brumell, Joanne Durgan, and Oliver Florey. 2022. “V-ATPase Is a Universal Regulator of LC3-Associated Phagocytosis and Non-Canonical Autophagy.” *The Journal of Cell Biology* 221 (6). https://doi.org/10.1083/jcb.202105112.
Ísaksson, Sigurjón. 2018. “Deeper Understanding of Autophagy and Pseudo-Autophagy through Latent-Space Analysis.” University of Cambridge. https://www.mlmi.eng.cam.ac.uk/files/dissertation-isaksson_reduced.pdf.
Iwatsuki-Horimoto, Kiyoko, Taisuke Horimoto, Takeshi Noda, Maki Kiso, Junko Maeda, Shinji Watanabe, Yukiko Muramoto, Ken Fujii, and Yoshihiro Kawaoka. 2006. “The Cytoplasmic Tail of the Influenza A Virus M2 Protein Plays a Role in Viral Assembly.” *Journal of Virology* 80 (11): 5233–40.
Jing, Li, Pascal Vincent, Yann LeCun, and Yuandong Tian. 2021. “Understanding Dimensional Collapse in Contrastive Self-Supervised Learning.” *arXiv [cs.CV]*. arXiv. http://arxiv.org/abs/2110.09348.
Johansen, Terje, and Trond Lamark. 2011. “Selective Autophagy Mediated by Autophagic Adapter Proteins.” *Autophagy* 7 (3): 279–96.
Karras, Tero, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2019. “Analyzing and Improving the Image Quality of StyleGAN.” *arXiv [cs.CV]*. arXiv.

http://arxiv.org/abs/1912.04958.
Kingma, Diederik P., and Max Welling. 2019. “An Introduction to Variational Autoencoders.” *Foundations and Trends® in Machine Learning* 12 (4): 307–92.
Klionsky, Daniel J., Amal Kamal Abdel-Aziz, Sara Abdelfatah, Mahmoud Abdellatif, Asghar Abdoli, Steffen Abel, Hagai Abeliovich, et al. 2021. “Guidelines for the Use and Interpretation of Assays for Monitoring Autophagy (4th edition)1.” *Autophagy* 17 (1): 1–382.
Kroemer, Guido, and Beth Levine. 2008. “Autophagic Cell Death: The Story of a Misnomer.” *Nature Reviews. Molecular Cell Biology* 9 (12): 1004–10.
Kügelgen, Julius von, Yash Sharma, Luigi Gresele, Wieland Brendel, Bernhard Schölkopf, Michel Besserve, and Francesco Locatello. 2021. “Self-Supervised Learning with Data Augmentations Provably Isolates Content from Style.” *arXiv [stat.ML]*. arXiv. http://arxiv.org/abs/2106.04619.
Kuo, Shu-Ming, Chi-Jene Chen, Shih-Cheng Chang, Tzu-Jou Liu, Yi-Hsiang Chen, Sheng-Yu Huang, and Shin-Ru Shih. 2017. “Inhibition of Avian Influenza A Virus Replication in Human Cells by Host Restriction Factor TUFM Is Correlated with Autophagy.” *mBio* 8 (3). https://doi.org/10.1128/mBio.00481-17.
Kuss-Duerkop, Sharon K., Juan Wang, Ignacio Mena, Kris White, Giorgi Metreveli, Ramanavelan Sakthivel, Miguel A. Mata, et al. 2017. “Influenza Virus Differentially Activates mTORC1 and mTORC2 Signaling to Maximize Late Stage Replication.” *PLoS Pathogens* 13 (9): e1006635.
Lamb, Christopher A., Tamotsu Yoshimori, and Sharon A. Tooze. 2013. “The Autophagosome: Origins Unknown, Biogenesis Complex.” *Nature Reviews. Molecular Cell Biology* 14 (12): 759–74.
Lam, Tommy Tsan-Yuk, Huachen Zhu, Jia Wang, David K. Smith, Edward C. Holmes, Robert G. Webster, Richard Webby, Joseph M. Peiris, and Yi Guan. 2011. “Reassortment Events among Swine Influenza A Viruses in China: Implications for the Origin of the 2009 Influenza Pandemic.” *Journal of Virology* 85 (19): 10279–85.
Le-Khac, Phuc H., Graham Healy, and Alan F. Smeaton. 2020. “Contrastive Representation Learning: A Framework and Review.” *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/2010.05113.
Li, Tian, Zhenyu Li, Erin E. Deans, Eva Mittler, Meisui Liu, Kartik Chandran, and Tijana Ivanovic. 2021. “The Shape of Pleomorphic Virions Determines Resistance to Cell-Entry Pressure.” *Nature Microbiology* 6 (5): 617–29.
Liu, Xiao, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. 2021. “Self-Supervised Learning: Generative or Contrastive.” *IEEE Transactions on Knowledge and Data Engineering*, 1–1.
Li, X. L., M. Boyanapalli, X. Weihua, D. V. Kalvakolanu, and B. A. Hassel. 1998. “Induction of Interferon Synthesis and Activation of Interferon-Stimulated Genes by Liposomal Transfection Reagents.” *Journal of Interferon & Cytokine Research: The Official Journal of the International Society for Interferon and Cytokine Research* 18 (11): 947–52.
Locatello, Francesco, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Jun 2019. “Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations.” In *Proceedings of the 36th International Conference on Machine Learning*, edited by Kamalika Chaudhuri and Ruslan Salakhutdinov, 97:4114–24. Proceedings of Machine Learning Research. PMLR.
Lowe, David G. 2004. “Distinctive Image Features from Scale-Invariant Keypoints.” *International Journal of Computer Vision* 60 (2): 91–110.
Ma, Jianhui, Qian Sun, Ruifang Mi, and Hongbing Zhang. 2011. “Avian Influenza A Virus H5N1 Causes Autophagy-Mediated Cell Death through Suppression of mTOR Signaling.” *Journal of Genetics and Genomics = Yi Chuan Xue Bao* 38 (11): 533–37.
Martin, Brigitte E., Jeremy D. Harris, Jiayi Sun, Katia Koelle, and Christopher B. Brooke. 2020. “Cellular Co-Infection Can Modulate the Efficiency of Influenza A Virus Production and

Shape the Interferon Response." *PLoS Pathogens* 16 (10): e1008974.
McEwan, David G. 2017. "Host-Pathogen Interactions and Subversion of Autophagy." *Essays in Biochemistry* 61 (6): 687–97.
McQuin, Claire, Allen Goodman, Vasiliy Chernyshev, Lee Kamentsky, Beth A. Cimini, Kyle W. Karhohs, Minh Doan, et al. 2018. "CellProfiler 3.0: Next-Generation Image Processing for Biology." *PLoS Biology* 16 (7): e2005970.
Menéndez, M. L., J. A. Pardo, L. Pardo, and M. C. Pardo. 1997. "The Jensen-Shannon Divergence." *Journal of the Franklin Institute* 334 (2): 307–18.
Miyato, Takeru, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. 2018. "Spectral Normalization for Generative Adversarial Networks." *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/1802.05957.
Ohsumi, Yoshinori. 2014. "Historical Landmarks of Autophagy Research." *Cell Research* 24 (1): 9–23.
Orvedahl, Anthony, and Beth Levine. 2008. "Viral Evasion of Autophagy." *Autophagy* 4 (3): 280–85.
Pan, Hong, Yijuan Zhang, Zichao Luo, Ping Li, Lanlan Liu, Ce Wang, Hanzhong Wang, Hongchang Li, and Yifan Ma. 2014. "Autophagy Mediates Avian Influenza H5N1 Pseudotyped Particle-Induced Lung Inflammation through NF-κB and p38 MAPK Signaling Pathways." *American Journal of Physiology. Lung Cellular and Molecular Physiology* 306 (2): L183–95.
Pappas, Claudia, Patricia V. Aguilar, Christopher F. Basler, Alicia Solórzano, Hui Zeng, Lucy A. Perrone, Peter Palese, Adolfo García-Sastre, Jacqueline M. Katz, and Terrence M. Tumpey. 2008. "Single Gene Reassortants Identify a Critical Role for PB1, HA, and NA in the High Virulence of the 1918 Pandemic Influenza Virus." *Proceedings of the National Academy of Sciences of the United States of America* 105 (8): 3064–69.
Parrish, Colin R., Pablo R. Murcia, and Edward C. Holmes. 2015. "Influenza Virus Reservoirs and Intermediate Hosts: Dogs, Horses, and New Possibilities for Influenza Virus Exposure of Humans." *Journal of Virology* 89 (6): 2990–94.
Perot, Brieuc P., Jeremy Boussier, Nader Yatim, Jeremy S. Rossman, Molly A. Ingersoll, and Matthew L. Albert. 2018. "Autophagy Diminishes the Early Interferon-β Response to Influenza A Virus Resulting in Differential Expression of Interferon-Stimulated Genes." *Cell Death & Disease* 9 (5): 539.
Piccinini, Filippo, Tamas Balassa, Abel Szkalisity, Csaba Molnar, Lassi Paavolainen, Kaisa Kujala, Krisztina Buzas, et al. 2017. "Advanced Cell Classifier: User-Friendly Machine-Learning-Based Software for Discovering Phenotypes in High-Content Imaging Data." *Cell Systems* 4 (6): 651–55.e5.
Radford, Alec, Luke Metz, and Soumith Chintala. 2015. "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks." *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/1511.06434.
Randall, Richard E., and Stephen Goodbourn. 2008. "Interferons and Viruses: An Interplay between Induction, Signalling, Antiviral Responses and Virus Countermeasures." *The Journal of General Virology* 89 (Pt 1): 1–47.
Ravikumar, Brinda, Marie Futter, Luca Jahreiss, Viktor I. Korolchuk, Maike Lichtenberg, Shouqing Luo, Dunecan C. O. Massey, et al. 2009. "Mammalian Macroautophagy at a Glance." *Journal of Cell Science* 122 (Pt 11): 1707–11.
Ren, Yizhong, Chufang Li, Liqiang Feng, Weiqi Pan, Liang Li, Qian Wang, Jiashun Li, et al. 2016. "Proton Channel Activity of Influenza A Virus Matrix Protein 2 Contributes to Autophagy Arrest." *Journal of Virology* 90 (1): 591–98.
Richemond, Pierre H., Jean-Bastien Grill, Florent Altché, Corentin Tallec, Florian Strub, Andrew Brock, Samuel Smith, et al. 2020. "BYOL Works Even without Batch Statistics." *arXiv [stat.ML]*. arXiv. http://arxiv.org/abs/2010.10241.

Rombach, Robin, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. “High-Resolution Image Synthesis with Latent Diffusion Models.” *arXiv [cs.CV]*. arXiv. http://arxiv.org/abs/2112.10752.
Rossman, Jeremy S., Xianghong Jing, George P. Leser, and Robert A. Lamb. 2010. “Influenza Virus M2 Protein Mediates ESCRT-Independent Membrane Scission.” *Cell* 142 (6): 902–13.
Rossman, Jeremy S., and Robert A. Lamb. 2011. “Influenza Virus Assembly and Budding.” *Virology* 411 (2): 229–36.
Roth, Kevin, Aurelien Lucchi, Sebastian Nowozin, and Thomas Hofmann. 2017. “Stabilizing Training of Generative Adversarial Networks through Regularization.” *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/1705.09367.
Schmid, D., and C. Münz. 2005. “Immune Surveillance of Intracellular Pathogens via Autophagy.” *Cell Death and Differentiation* 12 Suppl 2 (November): 1519–27.
Shaw, and Palese. 2007. “Orthomyxoviridae: The Viruses and Their Replication.” *Fields Virology*.
Shen, Shensi, Oliver Kepp, and Guido Kroemer. 2012. “The End of Autophagic Cell Death?” *Autophagy* 8 (1): 1–3.
Shorten, Connor, and Taghi M. Khoshgoftaar. 2019. “A Survey on Image Data Augmentation for Deep Learning.” *Journal of Big Data* 6 (1): 1–48.
Sims, Anna, Laura Burgess Tornaletti, Seema Jasim, Chiara Pirillo, Ryan Devlin, Jack Hirst, Colin Loney, et al. 2022. “Superinfection Exclusion Creates Spatially Distinct Influenza Virus Populations.” *bioRxiv*. https://doi.org/10.1101/2022.06.06.494939.
Sparrer, Konstantin M. J., Sebastian Gableske, Matthew A. Zurenski, Zachary M. Parker, Florian Full, Gavin J. Baumgart, Jiro Kato, et al. 2017. “TRIM23 Mediates Virus-Induced Autophagy via Activation of TBK1.” *Nature Microbiology* 2 (11): 1543–57.
Stevenson, J. P., and F. Biddle. 1966. “Pleomorphism of Influenza Virus Particles under the Electron Microscope.” *Nature* 212 (5062): 619–21.
Sun, Yang, Chenggang Li, Yuelong Shu, Xiangwu Ju, Zhen Zou, Hongliang Wang, Shuan Rao, et al. 2012. “Inhibition of Autophagy Ameliorates Acute Lung Injury Caused by Avian Influenza A H5N1 Infection.” *Science Signaling* 5 (212): ra16.
Tian, Yuandong, Xinlei Chen, and Surya Ganguli. Jul 2021. “Understanding Self-Supervised Learning Dynamics without Contrastive Pairs.” Edited by Marina Meila and Tong Zhang, Proceedings of Machine Learning Research, 139: 10268–78.
Tripathi, Shashank, Marie O. Pohl, Yingyao Zhou, Ariel Rodriguez-Frandsen, Guojun Wang, David A. Stein, Hong M. Moulton, et al. 2015. “Meta- and Orthogonal Integration of Influenza ‘OMICs’ Data Defines a Role for UBR4 in Virus Budding.” *Cell Host & Microbe* 18 (6): 723–35.
Ulferts, Rachel, Elena Marcassa, Lewis Timimi, Liam Changwoo Lee, Andrew Daley, Beatriz Montaner, Suzanne Dawn Turner, Oliver Florey, John Kenneth Baillie, and Rupert Beale. 2021. “Subtractive CRISPR Screen Identifies the ATG16L1/vacuolar ATPase Axis as Required for Non-Canonical LC3 Lipidation.” *Cell Reports* 37 (4): 109899.
Wait, Eric C., Michael A. Reiche, and Teng-Leong Chew. 2020. “Hypothesis-Driven Quantitative Fluorescence Microscopy - the Importance of Reverse-Thinking in Experimental Design.” *Journal of Cell Science* 133 (21). https://doi.org/10.1242/jcs.250027.
Wang, Kunfeng, Chao Gou, Yanjie Duan, Yilun Lin, Xinhu Zheng, and Fei-Yue Wang. 2017. “Generative Adversarial Networks: Introduction and Outlook.” *IEEE/CAA Journal of Automatica Sinica* 4 (4): 588–98.
Wang, Phil. 2020. *Byol-Pytorch*. Github. https://github.com/lucidrains/byol-pytorch.
Wang, Yingxue, Parul Sharma, Matthew Jefferson, Weijiao Zhang, Ben Bone, Anja Kipar, David Bitto, et al. 2021. “Non-Canonical Autophagy Functions of ATG16L1 in Epithelial Cells Limit Lethal Infection by Influenza A Virus.” *The EMBO Journal* 40 (6): e105543.
Wang, Yixin, and Michael I. Jordan. 2021. “Desiderata for Representation Learning: A Causal

Perspective.” *arXiv [stat.ML]*. arXiv. http://arxiv.org/abs/2109.03795.
Wan, Yushun, Wei Cao, Tao Han, Sheng Ren, Jian Feng, Tielong Chen, Jun Wang, Ruth Broering, Mengji Lu, and Ying Zhu. 2017. “Inducible Rubicon Facilitates Viral Replication by Antagonizing Interferon Production.” *Cellular & Molecular Immunology* 14 (7): 607–20.
Wileman, Tom. 2013. “Autophagy as a Defence against Intracellular Pathogens.” *Essays in Biochemistry* 55: 153–63.
Wise, Helen M., Edward C. Hutchinson, Brett W. Jagger, Amanda D. Stuart, Zi H. Kang, Nicole Robb, Louis M. Schwartzman, et al. 2012. “Identification of a Novel Splice Variant Form of the Influenza A Virus M2 Ion Channel with an Antigenically Distinct Ectodomain.” *PLoS Pathogens* 8 (11): e1002998.
Wu, Mike, Milan Mosse, Chengxu Zhuang, Daniel Yamins, and Noah Goodman. 2020. “Conditional Negative Sampling for Contrastive Learning of Visual Representations.” *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/2010.02037.
Wu, You-Tong, Hui-Ling Tan, Guanghou Shui, Chantal Bauvy, Qing Huang, Markus R. Wenk, Choon-Nam Ong, Patrice Codogno, and Han-Ming Shen. 2010. “Dual Role of 3-Methyladenine in Modulation of Autophagy via Different Temporal Patterns of Inhibition on Class I and III Phosphoinositide 3-Kinase *.” *The Journal of Biological Chemistry* 285 (14): 10850–61.
Xu, Yangqing, Junying Yuan, and Marta M. Lipinski. 2013. “Live Imaging and Single-Cell Analysis Reveal Differential Dynamics of Autophagy and Apoptosis.” *Autophagy* 9 (9): 1418–30.
Yanai, Itai, and Martin Lercher. 2020. “Selective Attention in Hypothesis-Driven Data Analysis.” *bioRxiv*. https://doi.org/10.1101/2020.07.30.228916.
Yan, Yiwu, Zhen Zou, Yang Sun, Xiao Li, Kai-Feng Xu, Yuquan Wei, Ningyi Jin, and Chengyu Jiang. 2013. “Anti-Malaria Drug Chloroquine Is Highly Effective in Treating Avian Influenza A H5N1 Virus Infection in an Animal Model.” *Cell Research* 23 (2): 300–302.
Yue, Zhenyu, Shengkan Jin, Chingwen Yang, Arnold J. Levine, and Nathaniel Heintz. 2003. “Beclin 1, an Autophagy Gene Essential for Early Embryonic Development, Is a Haploinsufficient Tumor Suppressor.” *Proceedings of the National Academy of Sciences of the United States of America* 100 (25): 15077–82.
Zang, Farong, Yinghu Chen, Zhendong Lin, Zhijian Cai, Lei Yu, Feng Xu, Jiaoli Wang, Weiguo Zhu, and Huoquan Lu. 2016. “Autophagy Is Involved in Regulating the Immune Response of Dendritic Cells to Influenza A (H1N1) pdm09 Infection.” *Immunology* 148 (1): 56–69.
Zhou, Zhi, Xuejun Jiang, Di Liu, Zheng Fan, Xudong Hu, Jinguo Yan, Ming Wang, and George F. Gao. 2009. “Autophagy Is Involved in Influenza A Virus Replication.” *Autophagy* 5 (3): 321–28.

5 - Author Contribution Statement

HCT116 ATG16L1 knockout cell lines and their derivatives, explored in section 2.2.1, were generated and validated by **Lewis Timimi**. The M2 constructs - and the pInd20 plasmids containing them - were generated by **Rachel Ulferts**. The data normalisation strategy used for processing of fluorescent images prior to their input to the deep learning models in section 2.1 was developed by **Sigurjon Isaksson**.

6 - Figures

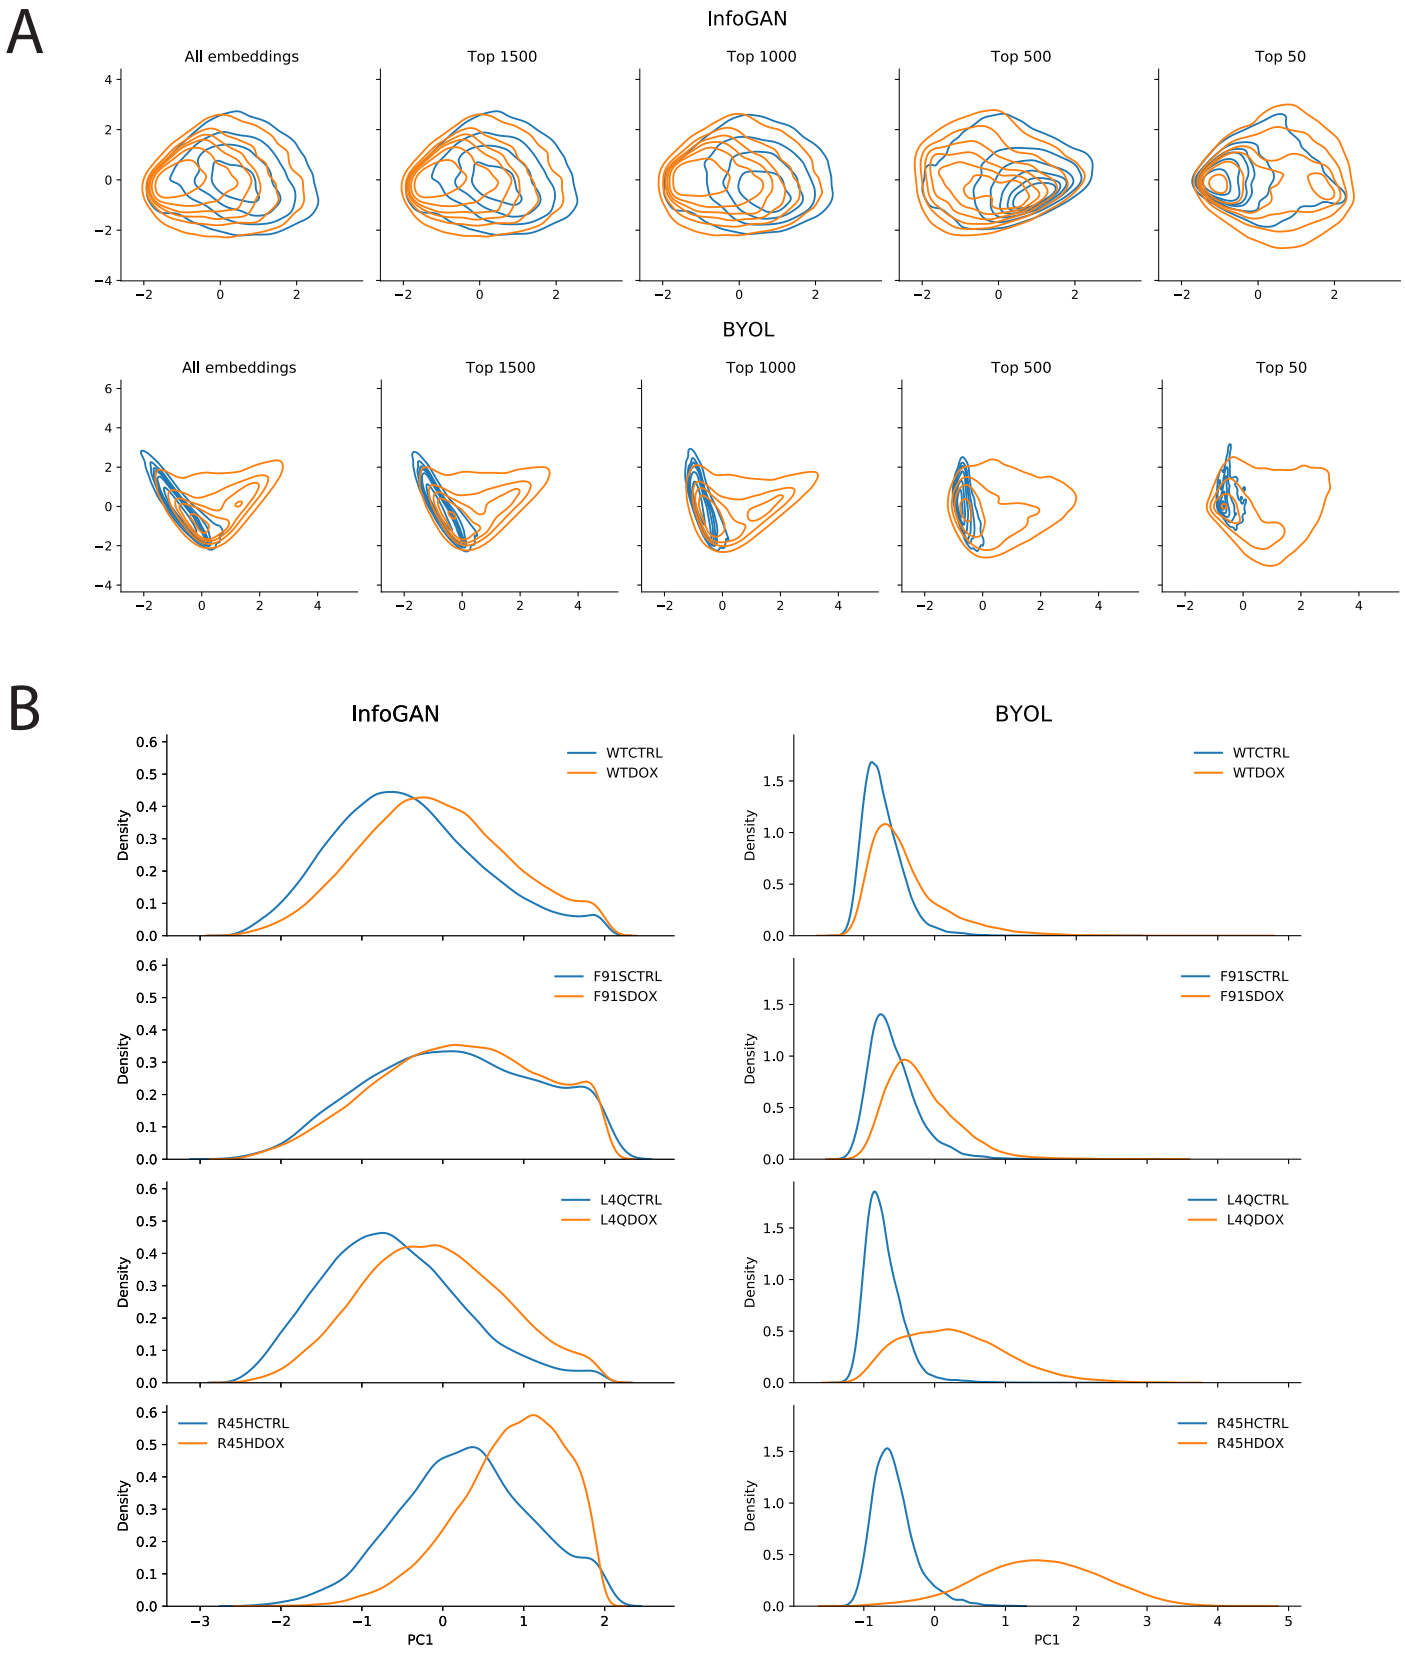

Figure 1: Image Representations generated by BYOL allow for better control of nuisance variation (A) PCA plots of InfoGAN- and BYOL-derived image representations progressively ablated of features that show low variance between control and treated conditions. (B) Kernel density estimation plots of distributions over the first principal component axis of variation of image feature embeddings. The embeddings were derived from the final convolutional layer of a custom InfoGAN discriminator network, and from the final maxpool layer from a WideResNet101 trained using BYOL. The most highly varying 1000 features were selected and fed into the PCA model. The x-axis of the InfoGAN distributions is inverted for easier visual comparison.

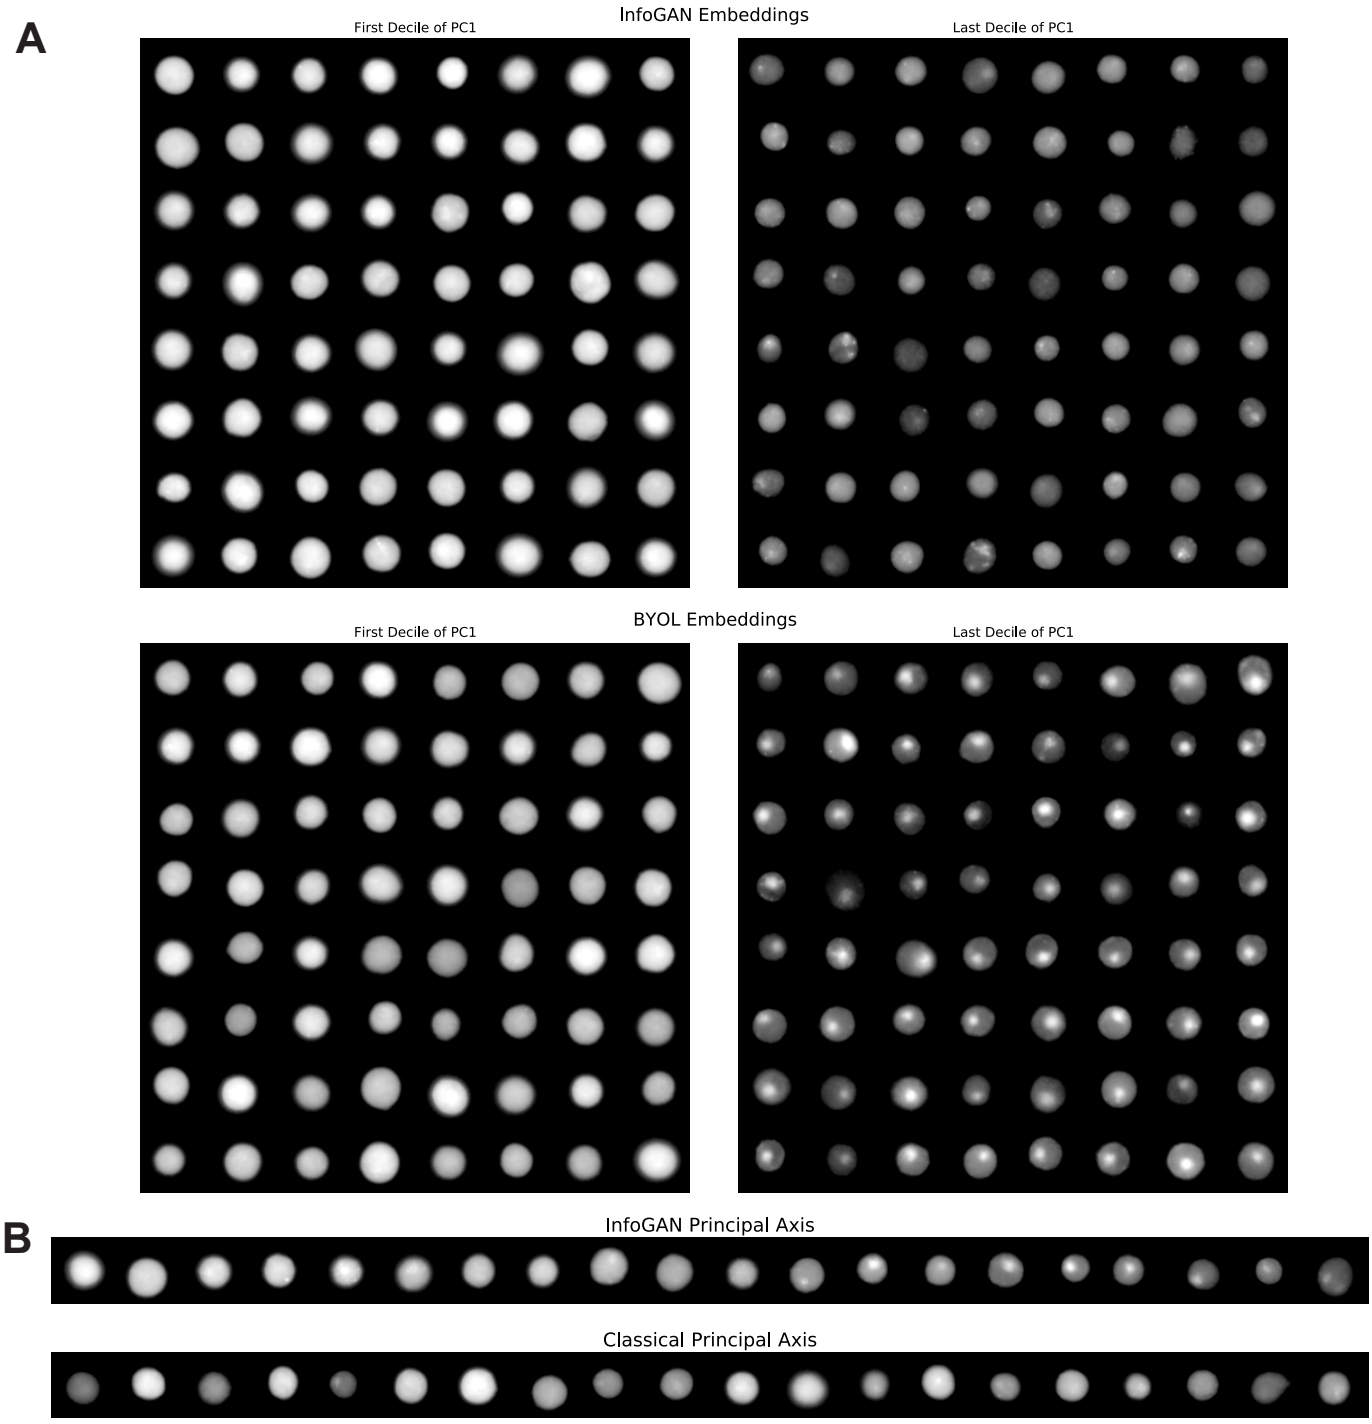

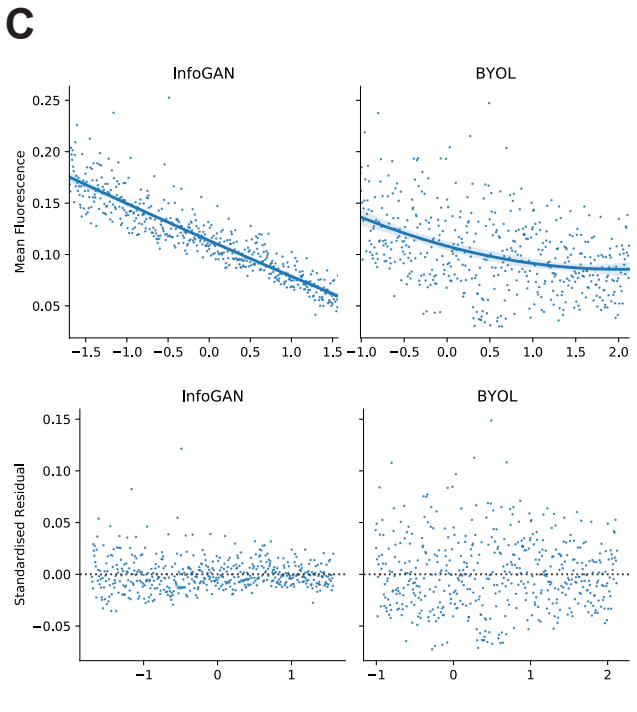

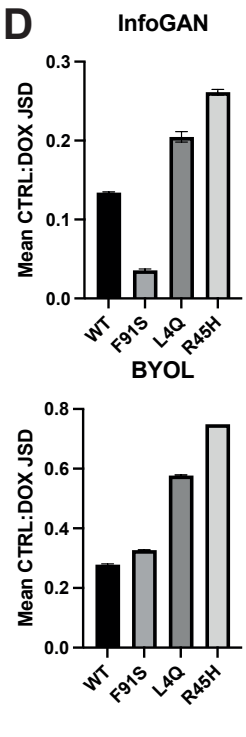

Figure 2: The primary axis of variation in BYOL embeddings corresponds to biological variation

(A) Cell images corresponding to the first and last deciles along the first principal component of InfoGAN- and BYOL-derived embeddings. (B) Images whose representations fall closest to a linear interpolation between the 5th and 95th percentile of the first principal component, PC1. (C) Regression and residual plots of mean fluorescence against the first principal component of BYOL and InfoGAN embeddings. Second order polynomial regression was performed, and bands show 95% CI. (D) Jensen-Shannon distance between CTRL and DOX populations along first principal component. Error bars show SD across five permutations of 750 features selected from top 1000 biologically varying features.

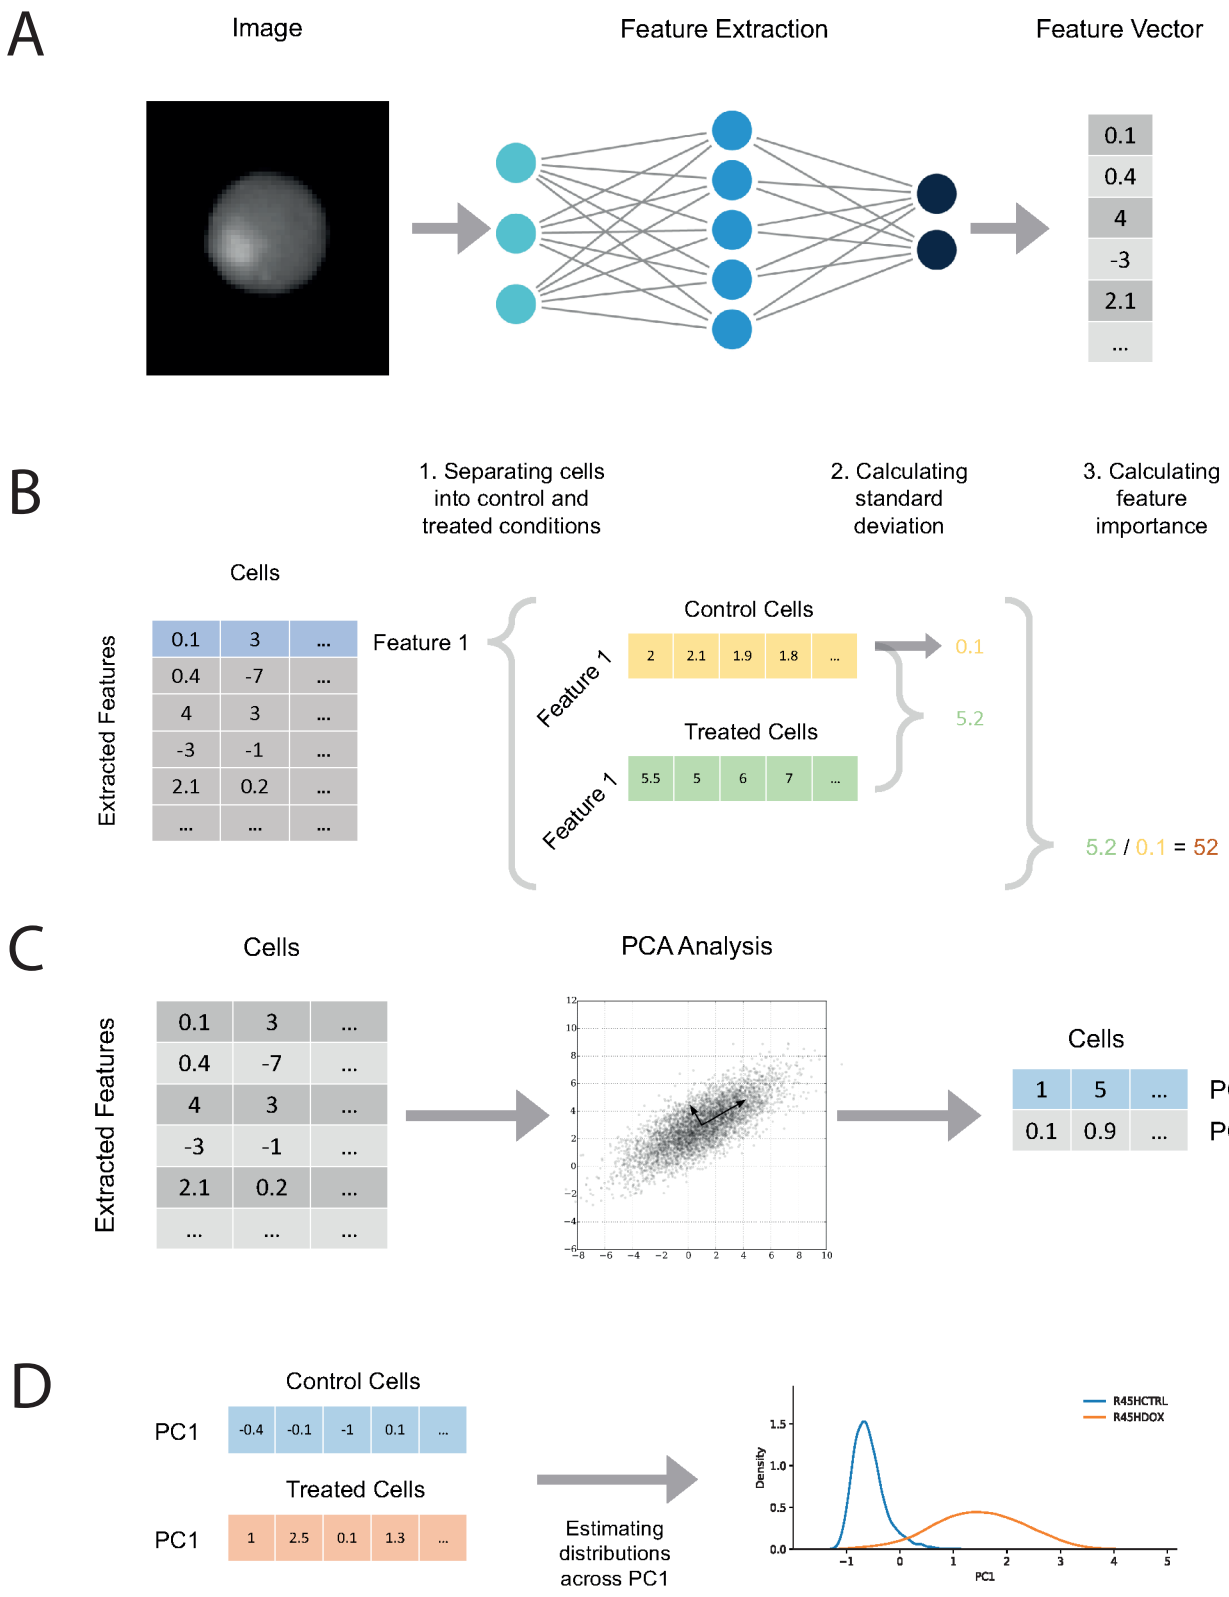

Figure 3: Feature values are processed by PCA, and their distributions along the first principal component give insight into the properties of the extracted features

(A) Images are processed by the trained neural network, generating a set of feature values. (B) Each cell image is then associated with a set of image features. Each feature dimension is selected, and separated into control and treated populations. The standard deviation is calculated across the control cells, and then both the control and treated cells. The ratio of these two measures of variance corresponds roughly to biological importance, allowing for feature selection. (C) After any selection, the resultant feature matrix undergoes two-component principal component analysis, generating a further matrix where each column is associated with a cell and the two rows correspond to where that particular cell lies along the first- and second-largest axes of variation in the dataset. (D) As the second principal component is poor at separating control and treated cells, it is discarded, and the distributions of cells in different conditions are plotted.

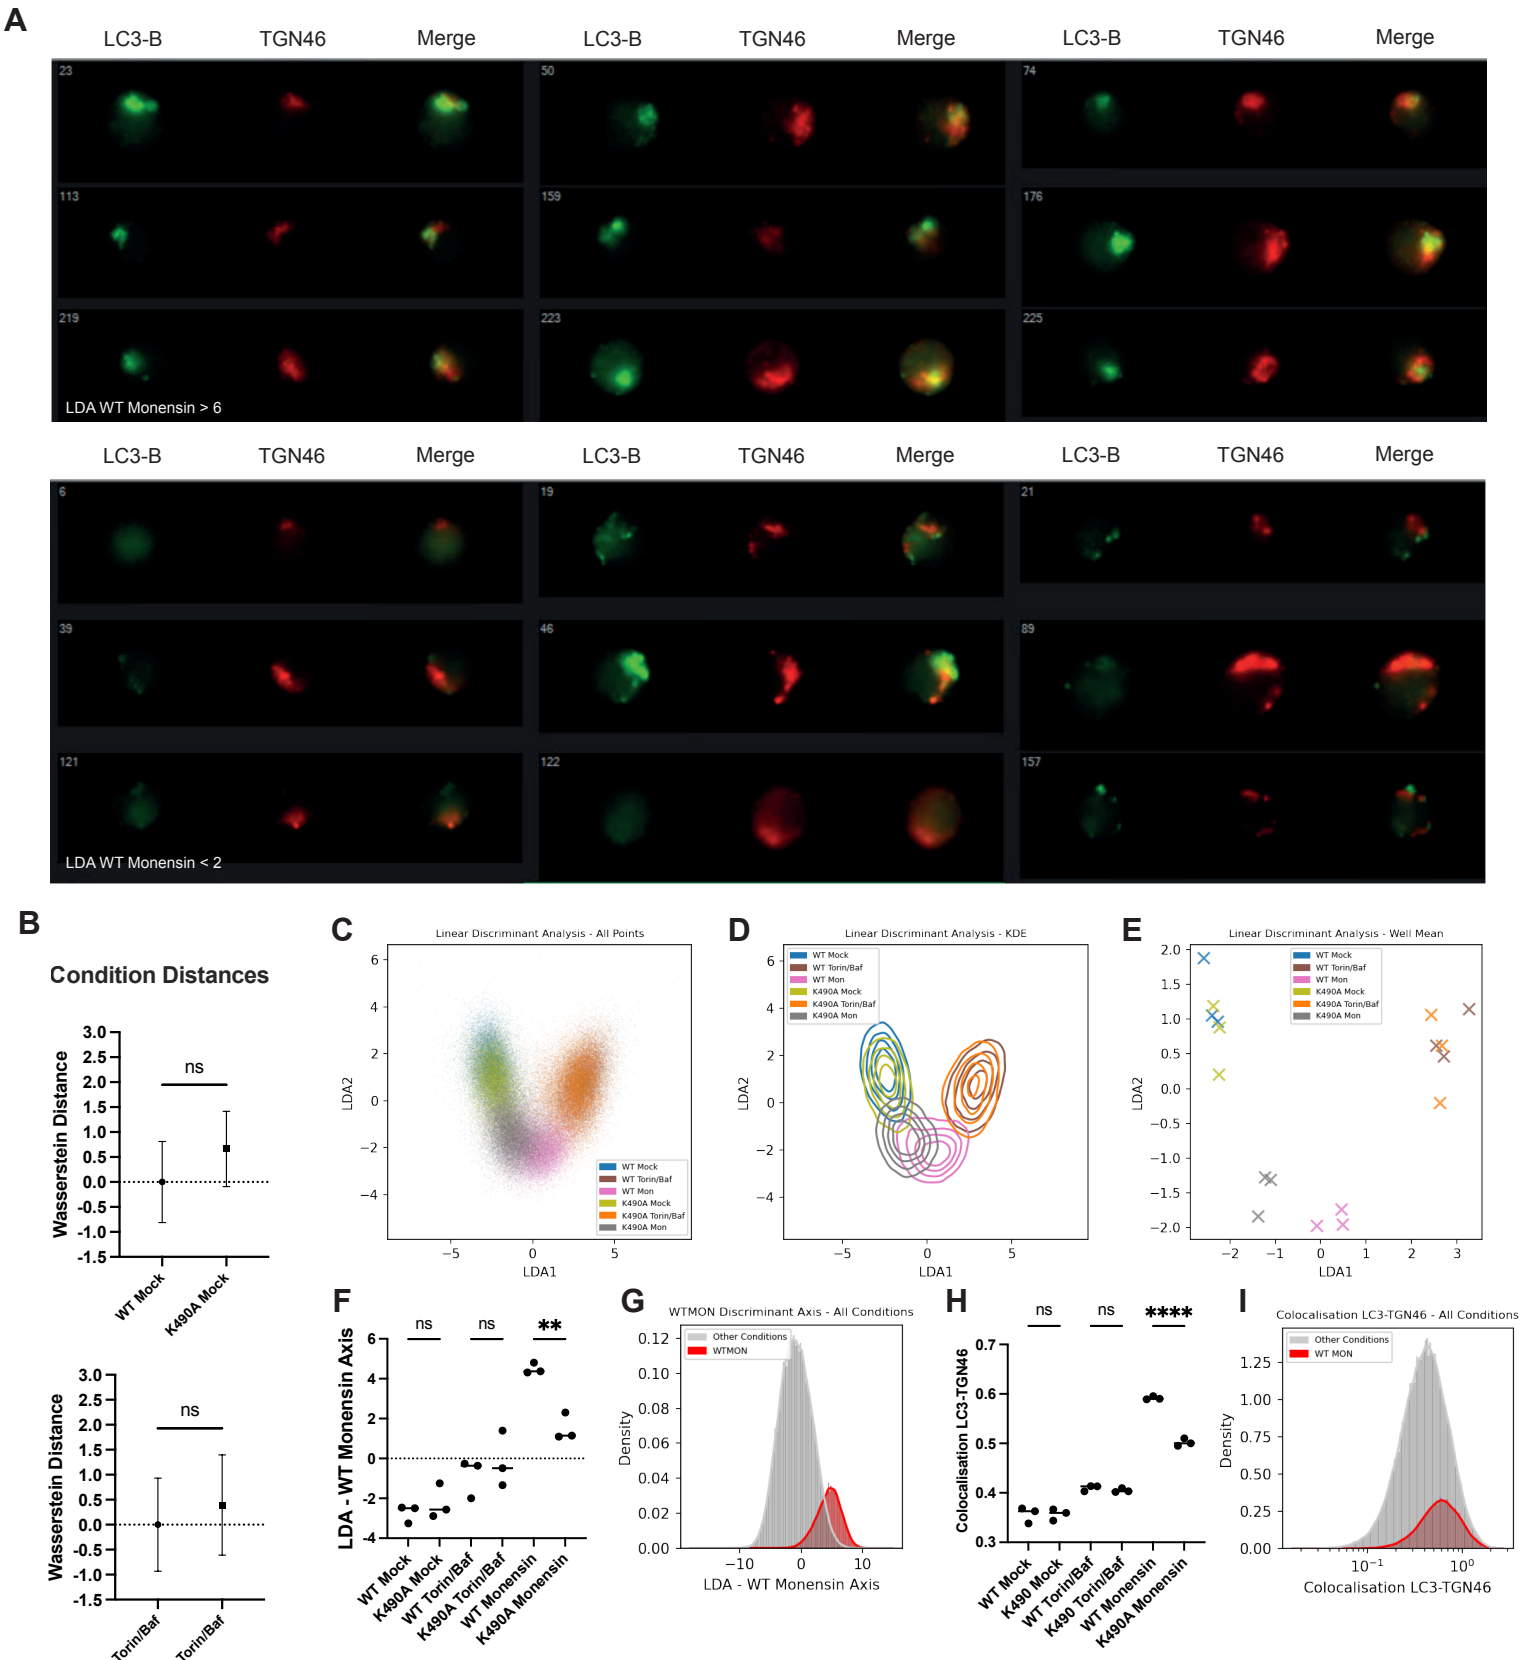

Figure 4: Distribution distance measures capture biological information in classically-extracted feature spaces A) Images taken from the WT Monensin population that possess a LDA WT Monensin axis value greater than 6 and less than 2. B) Wasserstein distances between conditions with distances between repeats within conditions taken as a standard deviation for statistical testing. C) Scatterplot of all datapoints after embedding into a two-dimensional linear discriminant analysis space. D) Kernel density estimate of all conditions in two-dimensional LDA space. E) Embedding means across repeats within conditions in two-dimensional LDA space. F) Feature data transformed along axis most discriminable for the WT Monensin phenotype, with one-way ANOVA statistical significance testing across genotypes for mean value of this phenotypic measurement. G) Histogram of values along the LDA WT Monensin axis. H) One-way ANOVA statistical significance testing across genotypes of Bright Detail Similarity between LC3-B and TGN46 image channels as a measure of colocalisation. I) Histogram of colocalisation values.

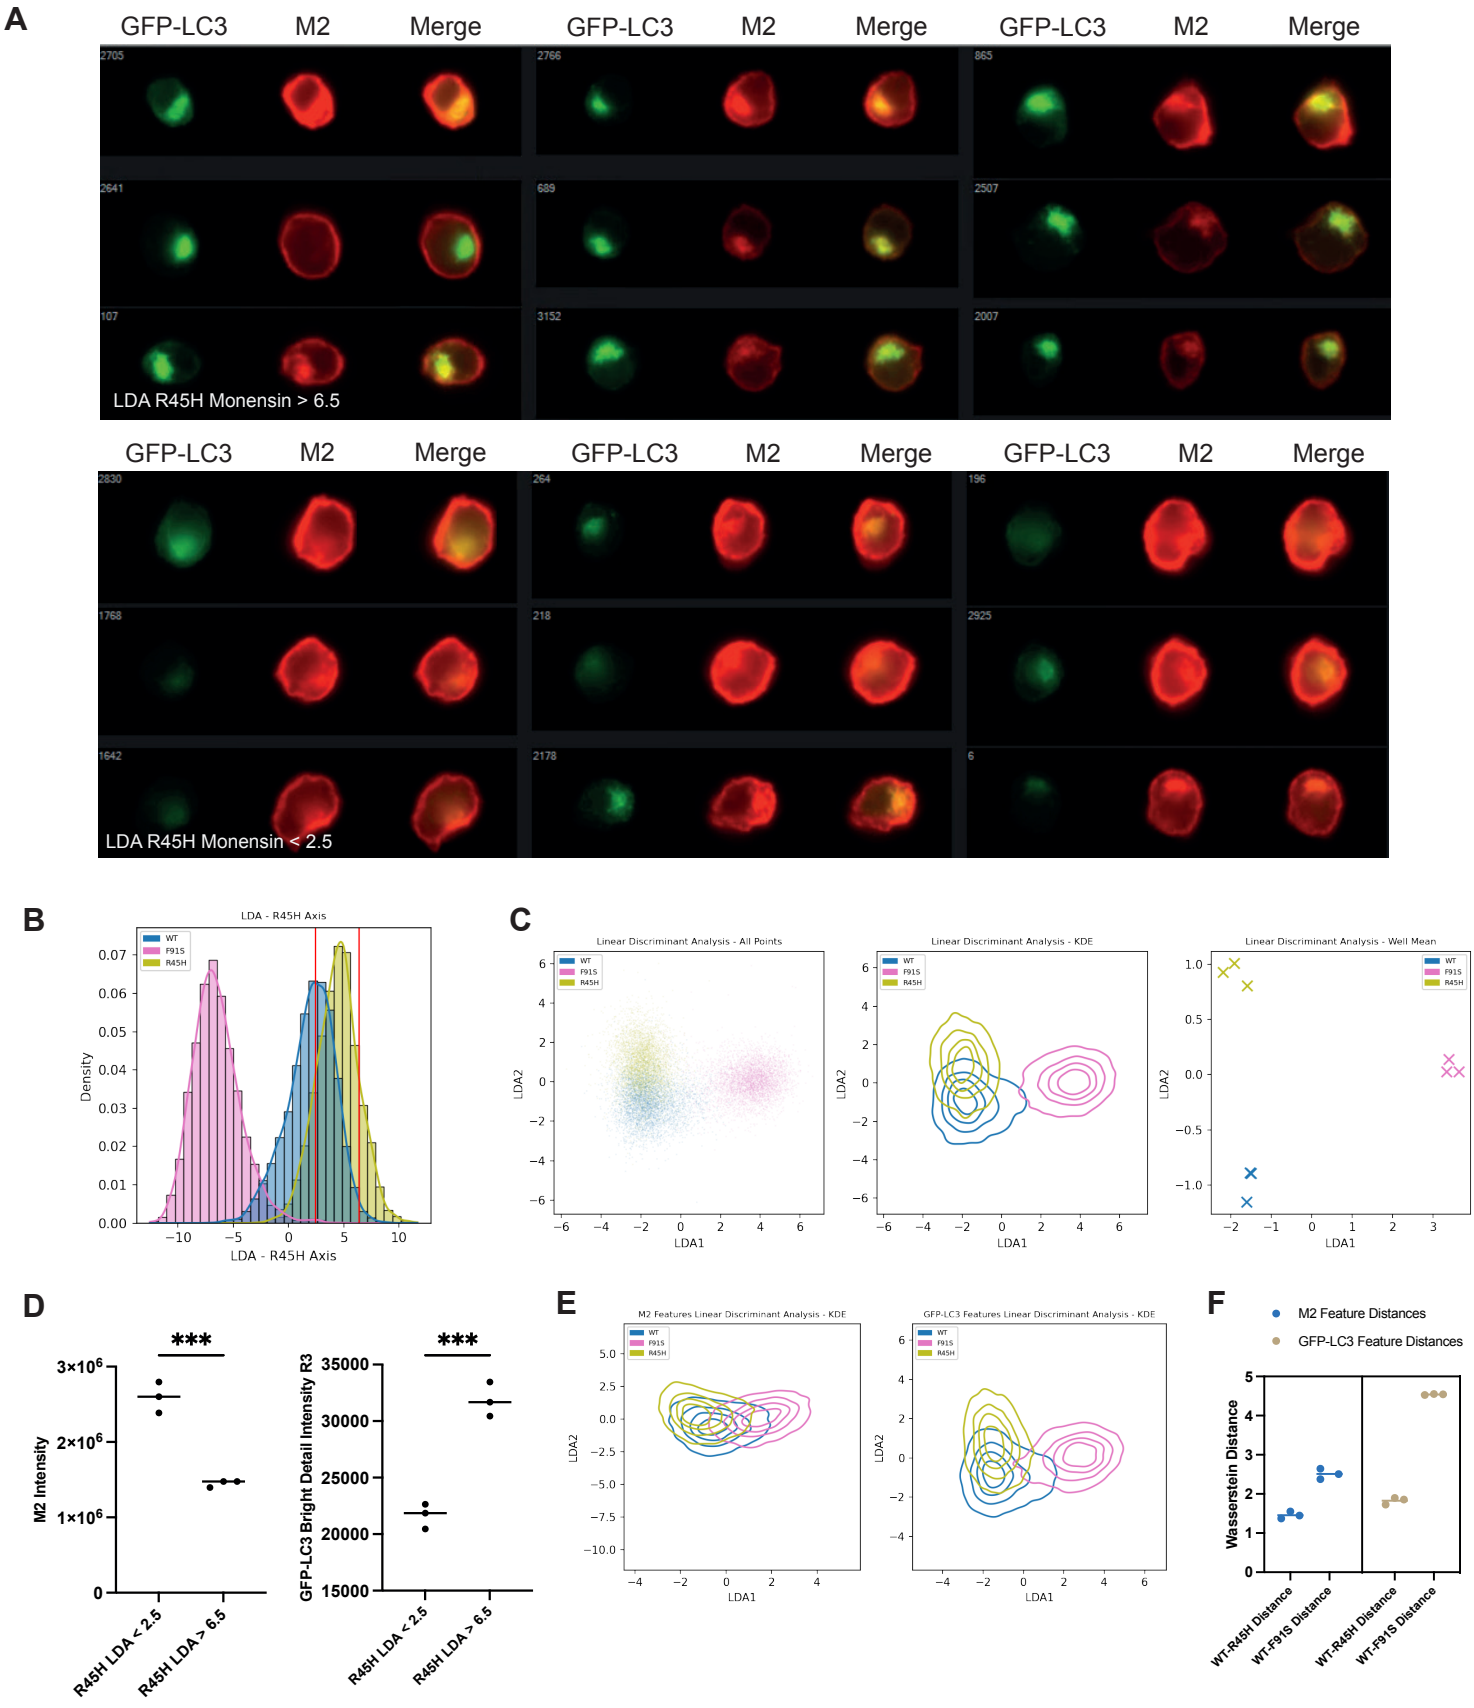

Figure 5: Distribution distance measures allow for phenotypic discovery in influenza M2 protein mutants
A) Representative fluorescence images of cells with LDA R45H values within the top and bottom 15th percentiles. B) Feature data transformed along axis most discriminable for the R45H phenotype, with top and bottom 15th percentiles of the R45H population displayed. C) Scatterplot and kernel density estimate of all datapoints after embedding into a two-dimensional linear discriminant analysis space, with embedding means across repeats. D) Quantification of average M2 fluorescence intensity and a measure of GFP-LC3 relocalisation between the most and least discriminable R45H populations, statistical significance testing performed per feature with standard t-test. E) Kernel density estimate of all conditions in two-dimensional LDA space generated using a) only features derived from the M2 fluorescence channel and b) only features derived from the GFP-LC3 fluorescence channel. F) Wasserstein distances between conditions derived from LDA performed on M2 and GFP-LC3 feature subsets.

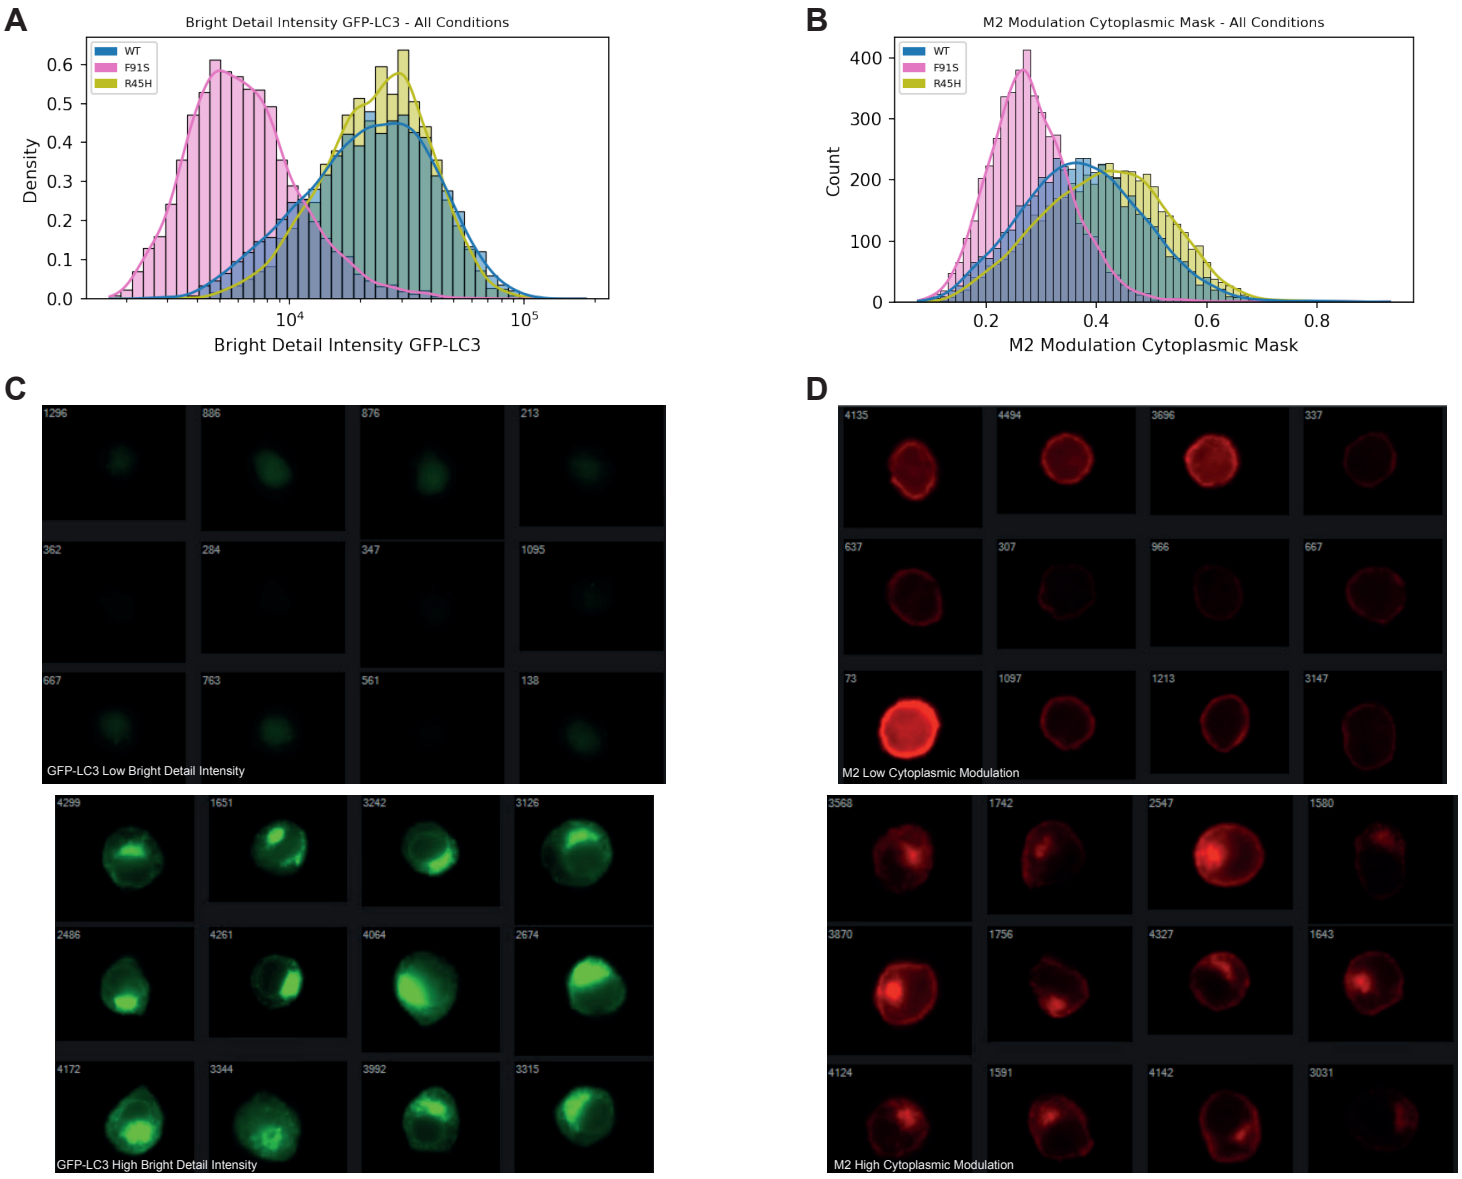

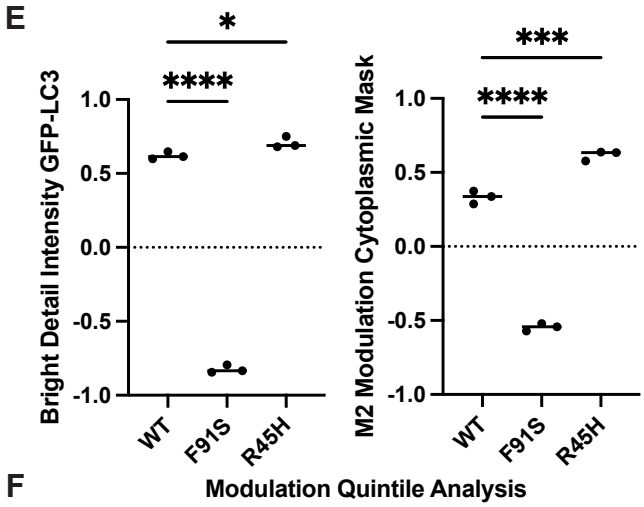

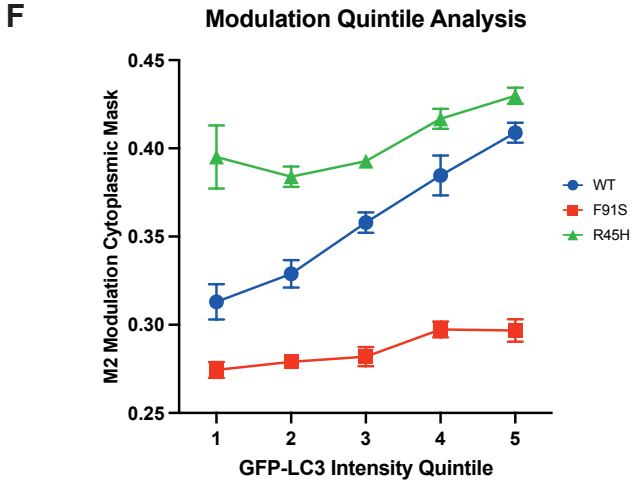

Figure 6: Different M2 Mutants display differences in GFP-LC3 relocalisation in addition to altered M2 localisation A) Histogram of the Bright Detail Intensity feature across all conditions, a constructed image feature that captures LC3 relocalisation. B) Histogram of intensity modulation within a cytoplasmic mask in the M2 channel. C) Representative images of cells with low and high Bright Detail Intensity values in the GFP-LC3 channel. D) Representative images of cells with low and high intensity modulation values in the M2 channel. E) Quantification of patterning differences across M2 mutants, statistical significance testing performed with one-way ANOVA. F) Quantification of intensity modulation within the M2 cytoplasmic mask across different quintiles GFP-LC3 intensity.

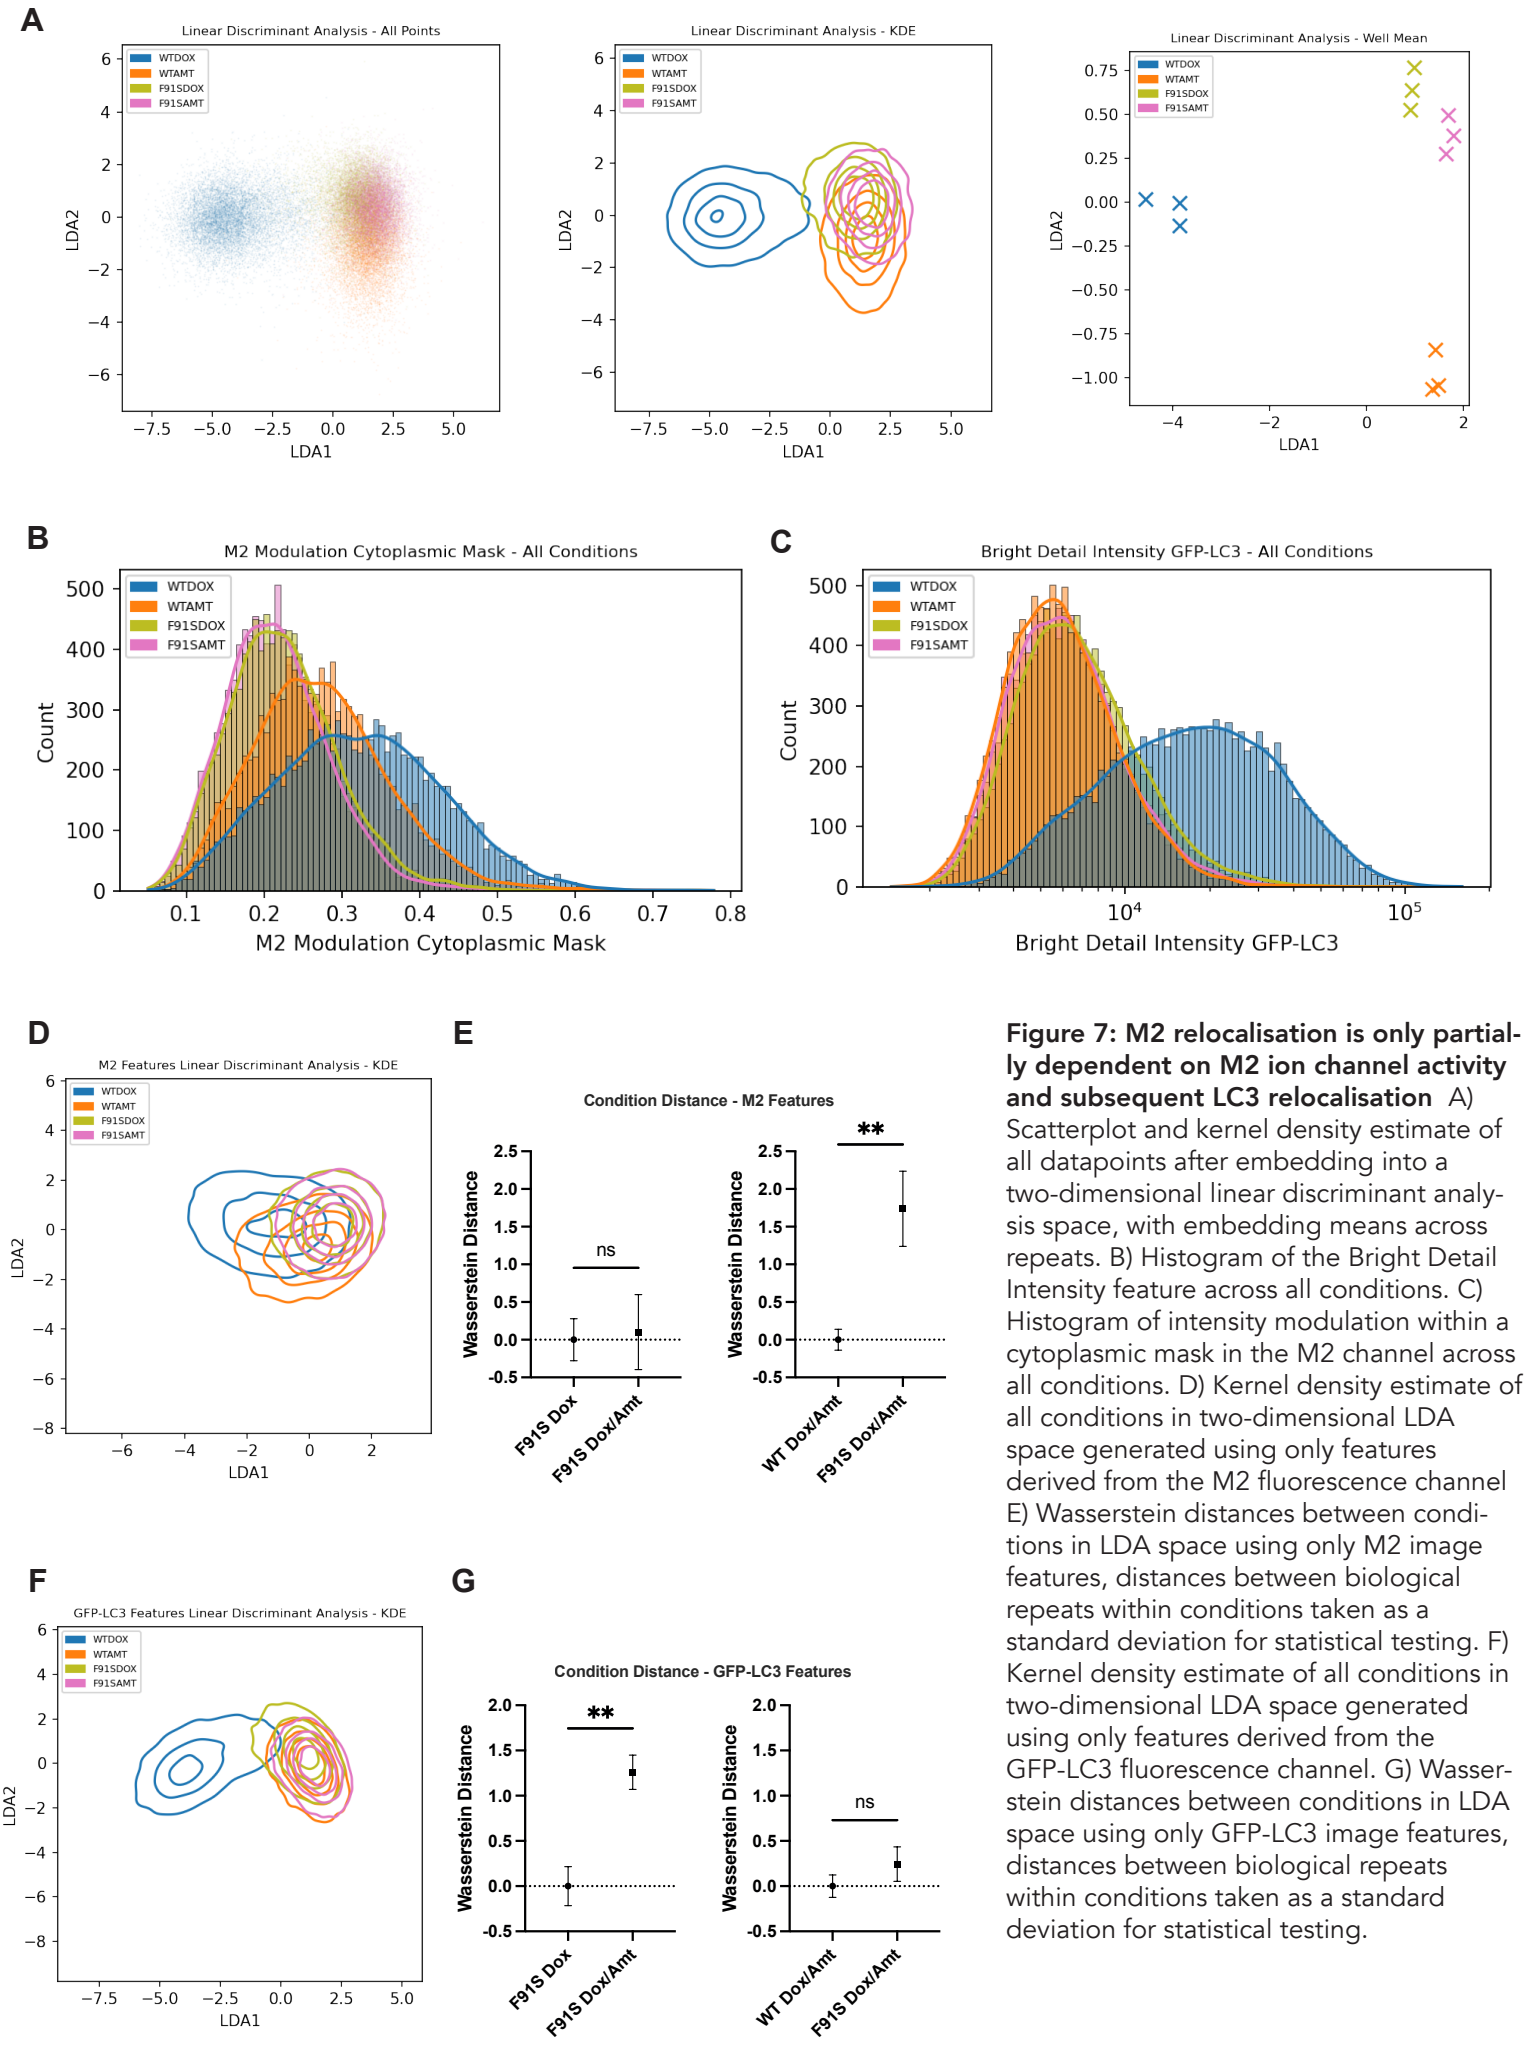

Figure 7: M2 relocalisation is only partially dependent on M2 ion channel activity and subsequent LC3 relocalisation A) Scatterplot and kernel density estimate of all datapoints after embedding into a two-dimensional linear discriminant analysis space, with embedding means across repeats. B) Histogram of the Bright Detail Intensity feature across all conditions. C) Histogram of intensity modulation within a cytoplasmic mask in the M2 channel across all conditions. D) Kernel density estimate of all conditions in two-dimensional LDA space generated using only features derived from the M2 fluorescence channel E) Wasserstein distances between conditions in LDA space using only M2 image features, distances between biological repeats within conditions taken as a standard deviation for statistical testing. F) Kernel density estimate of all conditions in two-dimensional LDA space generated using only features derived from the GFP-LC3 fluorescence channel. G) Wasserstein distances between conditions in LDA space using only GFP-LC3 image features, distances between biological repeats within conditions taken as a standard deviation for statistical testing.