AI-driven cell-fate prediction in microscopy

Rita Carlota[1,*], Mario Del Rosario[1,2,3,*], Inês Cunha[4], Juliette Griffié[4], Guillaume Jacquemet[5,✉], and Ricardo Henriques[1,6,✉]

[1]Instituto de Tecnologia Química e Biológica António Xavier, Universidade Nova de Lisboa, Oeiras, Portugal
[2]Faculty of Science and Engineering, Cell Biology, Åbo Akademi University, Turku, Finland
[3]InFLAMES Research Flagship Centre, University of Turku, Turku, Finland
[4]Science for Life Laboratory, Department of Biochemistry and Biophysics, Stockholm University, Stockholm, Sweden
[5]Faculty of Science and Engineering, Cell Biology, Åbo Akademi University; Turku Bioscience Centre, University of Turku and Åbo Akademi University, Turku, Finland
[6]Laboratory for Molecular Cell Biology, University College London, London, United Kingdom
[*]Equally contributed authors

Microscopy is evolving from a descriptive to a predictive tool in biology. Trained on time-lapse images, deep learning can forecast whether a cell will divide, die, or differentiate from how it looks and moves, in some systems, hours or even generations before the usual molecular markers appear, and without added fluorescent labels. Two questions frame this review and stay largely open: which image features actually carry the predictive signal, and whether a model foresees a genuinely future outcome or instead reads a state the cell has already entered. We organise the field by the visual signal that carries fate: changes in cell and nuclear shape over time, changes in brightness and texture, the dynamics of differentiation and competition, and patterns of movement. For each, we ask what must be measured, and over what area and time window, to predict fate. We then set out the main families of models, from networks that read a single image to those that read sequences of images over time and those that compress images into compact numerical summaries, outlining their strengths and limitations. Our central argument is that prediction has outrun validation: most reported performance is checked retrospectively against endpoint markers instead of on genuinely future cells. We close on the shared datasets, benchmarks, and interpretability work this gap demands, and on underexplored fates such as migration and senescence.

cell fate prediction | deep learning | live-cell imaging | morphodynamics | cell tracking | cell death | cell cycle | lineage tracing | computer vision | time-lapse microscopy
Correspondence: ***(G. Jacquemet) guillaume.jacquemet@utu.fi; (R. Henriques) r.henriques@itqb.unl.pt***

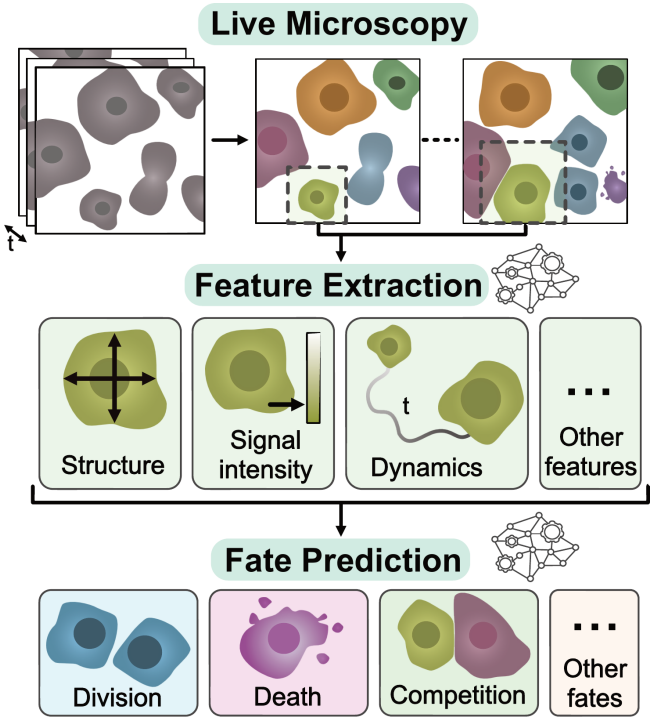

Fig. 1. Cell Fate Prediction Overview. Schematic representation of how different features extracted from microscopy data (e.g. morphology, signal intensity, dynamics) can be used, together with AI-driven models, to predict distinct cell fates.

Introduction

Cells combine order with randomness. Within a given cell type and environment, their behaviour is broadly reproducible: they grow, divide, differentiate or die in recognisable sequences. Individual cells still reach different outcomes, or reach the same outcome at different times, and this variability is usually called stochastic. At the level of single molecules or single snapshots, it can look genuinely random. Much of it instead reflects incomplete observation: structure that is invisible in one frame becomes legible once a cell is followed over time. Cellular behaviour is therefore better seen as patterned than as either fully deterministic or fully random, and this underlying structure is what makes fate prediction possible (1).

Cellular decision-making is seldom instantaneous; it emerges from cumulative biophysical and biochemical processes. As cells grow, divide, or differentiate, they undergo continuous changes in shape and organisation that reflect their internal regulatory state, shaped further by neighbouring cells and the local environment (2). If these physical changes link directly to molecular regulatory networks, then how cell shape, organisation, and signal intensity change over time may encode predictive information about future fate. Consistent with this, large-scale imaging that maps thousands of proteins in single cells shows that a cell's shape is tied to which proteins it expresses: two cells in the same cell-cycle phase but with different shapes carry distinct proteomic profiles that may already bias them toward different fates (3). Seen this way, phenotypic changes, including structural variations, intensity fluctuations, and temporal dynamics, may provide early indicators of fate decisions, although the precise limits of such predictability remain an open question.

A range of experimental techniques, including live-cell microscopy, immunofluorescence, flow cytometry, and quantitative image analysis, are widely used to read out cell state and cell-cycle progression (4). These approaches have revealed fundamental aspects of cellular regulation, yet many rely on molecular readouts that become detectable only after key regulatory changes have progressed. Many assays also capture static snapshots instead of continuous trajectories, lim-

iting access to the temporal information that may hold early predictive signatures (5).

Combining these biochemical approaches with artificial intelligence, especially deep learning, has widened the ability to extract predictive signals from such data and to turn them into predictions of future cell states. These methods identify features that are hard to quantify manually and integrate temporal information across long sequences. Cell Painting (6), for instance, stains several cellular components and organelles with a panel of fluorescent dyes imaged in separate colour channels. Software then outlines each cell and measures hundreds of numbers describing its size, shape, texture, and brightness. Together, these measurements pick up small differences between cells and increasingly link how a cell looks to what it is doing. Large reference datasets, such as the JUMP Cell Painting collection of around three million images across matched chemical and genetic perturbations, now underpin these image-based profiling efforts (7).

Neural networks trained on time-lapse microscopy data can predict aspects of cell fate, sometimes before conventional molecular markers become detectable (8). The reach and generality of this predictive capacity vary across systems and experimental conditions. Studies in stem-cell differentiation and cell-cycle progression, for example, show that time-resolved imaging improves predictive performance over static measurements. Complementary work indicates that imaging-derived features also reflect aspects of underlying molecular states (9, 10), supporting the idea that such features encode biologically relevant information about cell state and morphology (11–14).

Phenotypic changes, read over time and in their spatial context, can therefore serve as a functional readout of cellular dynamics. The field has diversified rapidly, with approaches spanning different cell types (15), fate outcomes, and modelling strategies. Some methods explicitly track lineages through mitosis, while others infer differentiation potential or reconstruct trajectories from limited temporal data. As imaging and molecular profiling become more integrated, morphology now sits alongside omics data within a wider, high-dimensional description of cell state.

Here we organise microscopy and AI-based cell-fate prediction by the dominant source of predictive signal: structural morphodynamics, optical and intensity changes, differentiation and competition dynamics, and motility and trajectories (Fig. 1). Within each, we ask a single question: what information, measured over which spatial and temporal scales, suffices to predict fate, and with which model family. Read across the field, a consistent pattern frames our argument: models increasingly forecast fate, often before molecular markers appear, yet seldom explain which features carry the signal or test their predictions prospectively. We use this gap to separate what is established from what remains only plausible, and to set out what the field would need to standardise to close it.

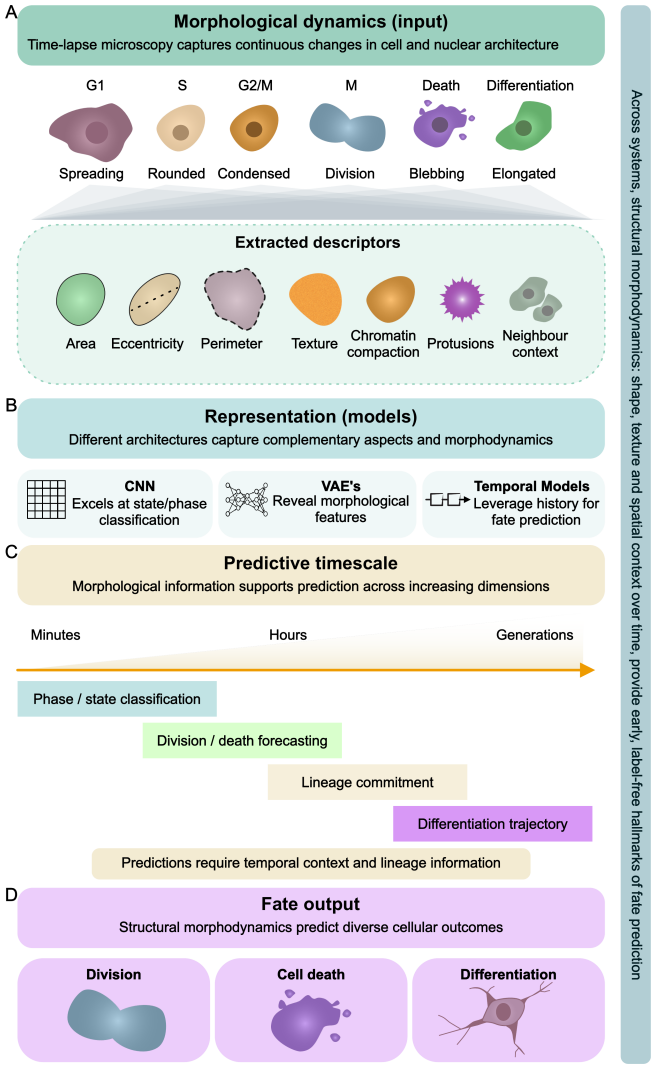

Fig. 2. Structural morphodynamics as predictors of cell fate. A. Time-lapse microscopy captures continuous changes in cell and nuclear architecture, including size, shape, texture, compaction, protrusions, and spatial context. B. These features can be analysed in different ways: networks that classify a cell's state from a single image, models that learn features on their own, and models that use a cell's history over time to predict its fate. C. Across biological contexts, structural morphodynamics support prediction of division, death, arrest, and differentiation before conventional endpoint markers become detectable. D. Predictions built on structural morphodynamical readouts can foresee a range of cellular outcomes.

Predicting Fate from Structural Morphodynamics

The cell cycle provides a natural test bed for morphology-based fate prediction. Its signature structural transitions, nuclear condensation, cell rounding, division and re-spreading, occur in reproducible sequences that time-lapse microscopy can capture (Fig. 2). Beyond simply assigning a cell to a phase, the key question is whether these shape changes can *predict future outcomes*, and how early. Convolutional neural networks (CNNs), models that learn visual patterns straight from image pixels, can read cell-cycle stage from a single fluorescence image (16) using features such as nuclear size and chromatin texture. They suit this task because they recognise a feature wherever it appears in the image and build it up in layers, from edges and textures to larger structures such

as condensed chromatin or a rounded cell. Their reach stays limited, though, when the history of the cell is ignored. Different fates can look almost identical in a single frame, so a one-frame model cannot reliably tell apart outcomes such as division and cell death. Fate-relevant information therefore lies in shape and in how that shape changes over time. Reassuringly, when researchers asked which features these models used, the models had rediscovered known criteria such as nuclear size and DNA content on their own, instead of being handed them. Other models, called β-variational autoencoders (β-VAEs), go further by compressing each cell image into a handful of numbers (a 'latent space') that vary smoothly with cell-cycle stage, replacing rigid phase labels with a continuous description. Together, these studies established cell shape as a self-contained, predictive readout of cell-cycle state.

Models that follow cells over time, by contrast, can predict outcomes early. Tracking how a cell's shape changes lets differentiation and fate decisions be called *hours before* the usual molecular markers appear. In the differentiating *Xenopus laevis* mucociliary epithelium (17), for example, a model trained on how single cells change over time predicted basal, goblet and multiciliated fates with moderate but stable accuracy, and a cell's position and shape over time carried more information than any single snapshot. Cell shape alone likewise separated fat-forming from bone-forming differentiation in human stromal cells, with usable accuracy from early timepoints (18). Prediction improves again when the full three-dimensional shape is used instead of a flattened 2D view. Geometric deep-learning methods, which treat a cell as a 3D object instead of a flat image, such as MorphoMIL (19), pick up subtle shape features that a 2D projection loses, making it easier to tell which cytoskeletal or chemical treatment a cell has experienced. Related generative models learn the whole-cell 3D layout, capturing how organelles are arranged relative to the cell and its nucleus (20).

A recurring worry with models that compress images into a few numbers is knowing what those numbers mean, and AI4CellFate (21) is a helpful worked example. It reduces each single-cell image to just two numbers that, early during growth-inhibitory therapy, already separate cells that will keep proliferating from those that will not. When the authors varied each of the two numbers and watched how the reconstructed cell image changed, one turned out to track brightness and the other cell size, the two features that together drive the prediction. Tying the model's internal numbers back to visible cell properties turns an otherwise opaque model into a readout a biologist can interpret.

A clear progression appears. A single image is enough to tell a cell's current state; compressing images into a few numbers captures its features more flexibly; and following shape over time, ideally in 3D, is what enables *early prediction of fate*. In practice, division, death or differentiation can be called several hours before conventional markers, making shape dynamics an early, label-free predictor of what a cell will do.

Predicting Fate from Optical and Intensity Changes

How bright a cell appears, measured by fluorescence, by ordinary transmitted light, or by phase-contrast imaging, gives a second route to predicting fate. These brightness signals shift as cells grow, copy their DNA, and rearrange chromatin, leaving optical fingerprints that track both cell-cycle progression and commitment to a particular fate (Fig. 3).

Early work showed that brightness alone can sort cells into cell-cycle phases from single images. Quantitative phase imaging (QPI), for instance, measures how much a cell slows the light passing through it, which reflects how much material the cell contains; this quantity rises steadily as the cell cycle advances. From this signal alone, models can tell G1, S, and G2/M apart without fluorescence and without any time information (22), showing that cell-cycle stage is written into a cell's optical properties.

Other models learn to do this without being told the phases in advance. A Cell Cycle Variational Auto-Encoder (CC-VAE) (23), for example, arranged fluorescent DNA-marker images along a continuous progression that replicates the cell cycle, using the natural rise in DNA signal during S phase and G2/M. Earlier studies had already shown that deep learning can reconstruct a smooth cell-cycle path from single-cell images with no time information (24), recover the cyclic order of the cell cycle straight from raw images (25), and place fixed cells in cell-cycle order to map how protein levels change through the cycle, all without any cell-cycle reporter (26). Together, these approaches show that brightness and image features can be turned into an ordered description of where a cell sits in the cycle.

As with shape, prediction improves once change over time is included. Time-lapse imaging of glioblastoma stem-like cells under ordinary transmitted light (27) shows that a model reading brightfield images over time can separate dividing from dying cells. Label-free models can even sort several outcomes at once, including division, apoptosis, necrosis and senescence, from phase-contrast images (28). With no fluorescence at all, small changes in transmitted-light brightness and contrast, tracked over time, carry enough information to predict what a cell will do.

These label-free models carry a risk, though. They may latch onto quirks of one dataset instead of real biology, and the way time is built into them can smear out rare cell types or sudden changes. CC-VAE, for instance, is trained so that successive frames of the same cell stay close together in its numerical summary, which keeps trajectories smooth over time. Whether forcing neighbouring frames to stay close also dulls the model's sensitivity to sudden fate switches is an open concern: a cell hit by DNA damage or acute stress may change abruptly in a way that a smoothness-favouring model glosses over. This tension between real, abrupt biological transitions and models built to prefer smooth ones is unresolved.

The pattern resembles the shape-based case: a single brightness measurement classifies cell-cycle phase, a compact nu-

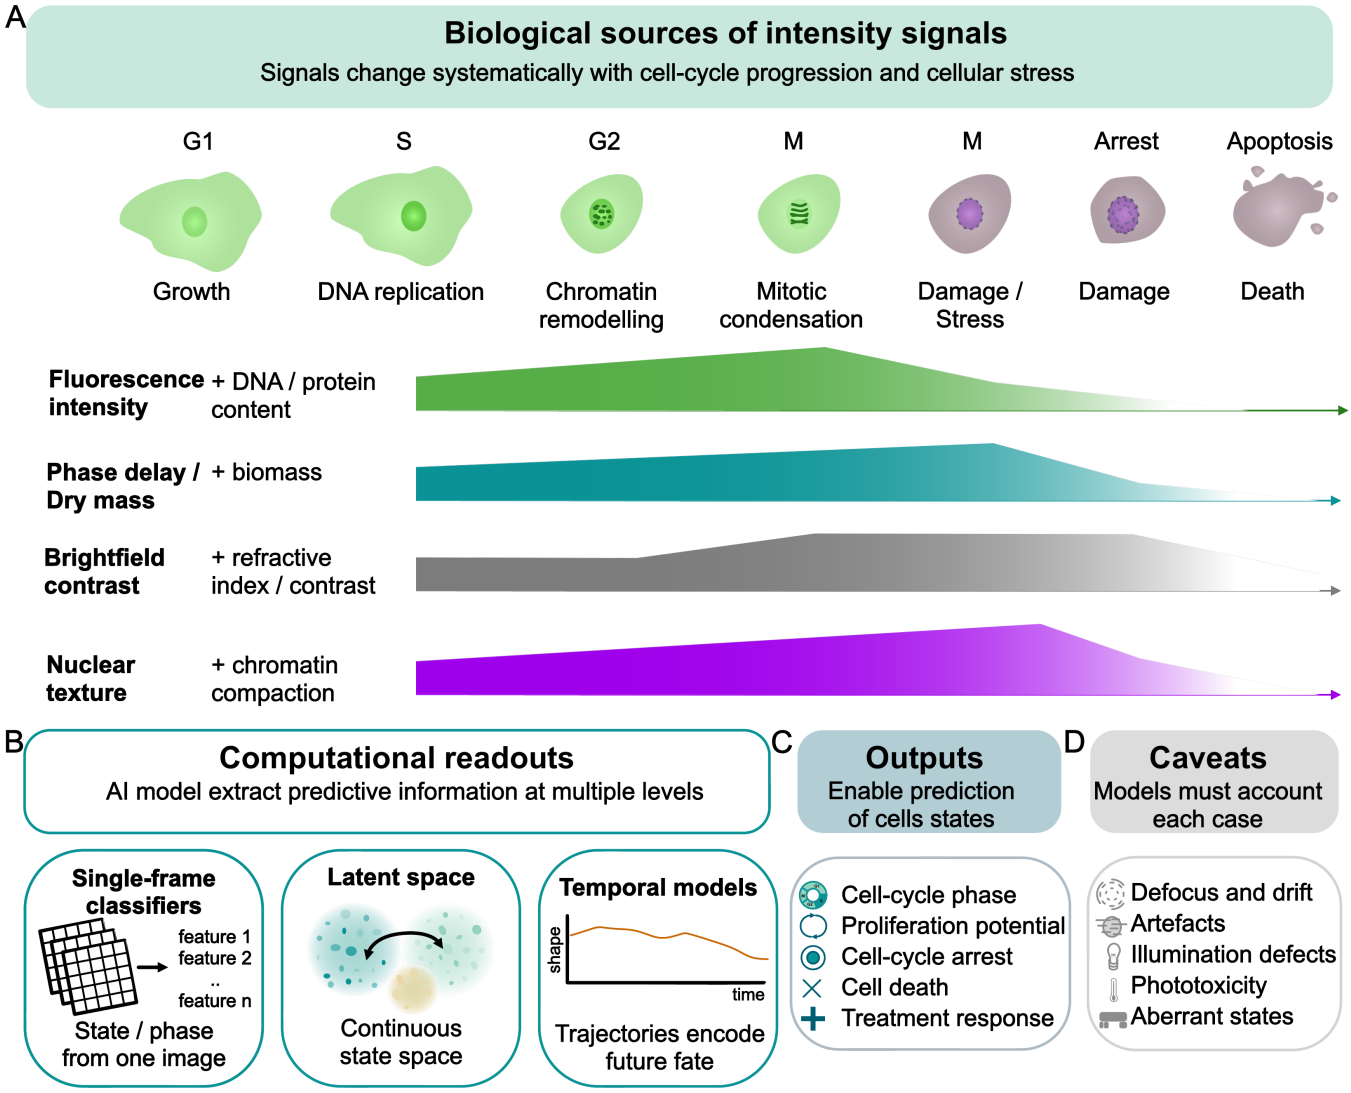

Fig. 3. Optical intensity trajectories as label-free proxies of internal cell state. A. Changes in fluorescence intensity, phase delay, brightfield contrast, and texture reflect underlying biomass accumulation, DNA replication, chromatin remodelling, and stress-associated reorganisation. B. AI models can use these signals in three ways: classifying single frames directly, compressing the optical signal into a compact numerical summary, and modelling how brightness changes over time. C and D. These approaches predict cell-cycle phase and fate-related outcomes without added fluorescent labels, but remain sensitive to imaging artefacts, uneven illumination, and abrupt state changes.

merical summary captures progression, and following brightness over time supports early prediction of fate. None of this needs added molecular markers; it rides on the natural optical changes cells undergo, which makes brightness-based imaging a gentle, scalable route to both a cell's state and its future behaviour.

Predicting Fate from Differentiation and Competition Dynamics

Beyond the shape and brightness changes seen across the cell cycle, many cells change in characteristic ways as they differentiate. As a cell moves along a developmental path, it passes through different appearances and short-lived intermediate states (Fig. 4). During the differentiation of induced pluripotent stem cells (iPSCs), for example, cells with visibly different shapes appear, so how a cell looks can already hint at where it is headed.

One deep learning framework, DEEP-MAP (29), was built to profile the shape and behaviour of human pluripotent stem cells (hPSCs). The cells carried two fluorescent markers: a labelled nucleus and a two-colour FUCCI reporter that shows a cell's cell-cycle phase by its colour, so nuclear shape, proliferation, and cell-cycle stage could be watched at once. Combining these with proliferation measures, the authors mapped how cell states changed over several days and predicted the fate of single cells. Local cell density and being in G1 were the most informative features for telling cell states apart, with nuclear area and how cells moved also contributing. The appeal is that this read out fate in real time, cell by cell, using generic proliferation markers already common in many labs instead of a special reporter tailored to each lineage.

Tolonen and colleagues took a different route (17). They stained the tissue at the end to reveal each cell's final identity, then traced every cell backwards through the time-lapse to label its earlier frames, using standard tools to outline cells and follow them over time. Their model system was the

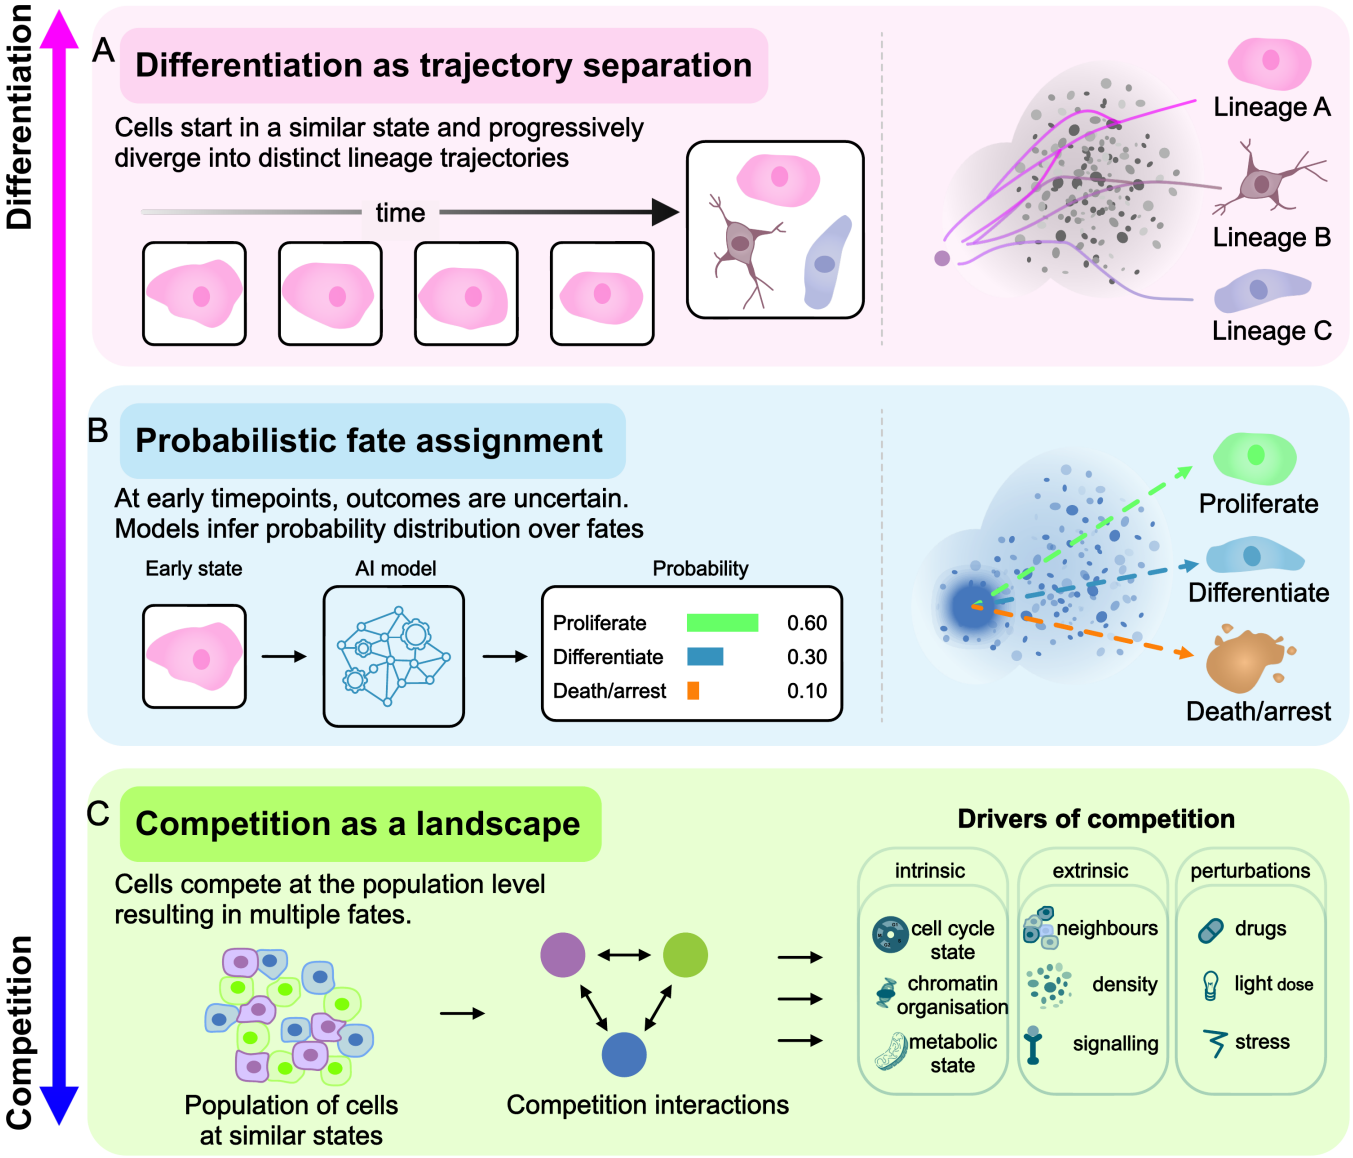

Fig. 4. Cell fate prediction through differentiation and competition dynamics. A. A cell's fate can be pictured as a path through a space of possible shapes and states over time. During differentiation, cells that start out similar gradually follow separate paths that become easier to tell apart as time goes on. B. When several fates compete, such as proliferation, differentiation, and death, a cell in an overlapping region can be assigned each outcome with some probability. C. Across a population, these outcomes reflect both the cell's own state and its surroundings, with cells pushed toward particular fates as they compete. This view links feature-based, probability-based, and path-based models of cell fate.

Xenopus laevis mucociliary epithelium, a fast-differentiating, easily imaged tissue containing several specialised cell types. Following each cell over time, they recorded how its shape, nuclear envelope, position, and movement changed as it went from an uncommitted progenitor to a specialised cell. Cell shape and position together carried substantial information about a cell's eventual fate, while movement contributed little, showing that shape dynamics alone can partly predict which lineage a cell will choose, even inside a crowded, living tissue.

Another model, a temporal-convolution-based variational autoencoder, τ-VAE, learned on its own to combine a cell's shape, how it changed over time, and how its neighbours were arranged, to predict the outcome of cell competition (2). It found local crowding to be the strongest predictor of fate during mechanical competition, and it let the authors read off which physical features and timescales mattered for an accurate call. By adding a component trained to spot when the outcome had shifted, they could also flag drugs that alter competition, showing the model could be used to screen candidate treatments.

Across these studies, features that change over time, such as cell shape and position, stand out as strong indicators of fate.

Predicting Fate from Motility and Trajectories

How a cell moves carries its own clues about fate, on top of what happens inside it. A cell's track over time records how it responds to its surroundings, including where it has been, how it interacts with neighbours, and the physical forces on it.

One model that reads both a cell's appearance and its movement over time, for example, was trained on small brightfield image crops plus how far each cell had moved, and predicted

whether blood progenitor cells would commit to one lineage or another, doing so before the usual fluorescent lineage markers could be seen (8). How the cell moved, together with the sequence of image crops, was enough to sort the cells, showing that lineage choice is already written into a cell's early movement and shape, visible under plain brightfield light. Other tools find such patterns without any labels: cellPLATO groups cells by their shape and movement to pick out distinct migrating subpopulations (30). Movement-based readouts extend naturally to whole tissues, where a cell's position relative to its neighbours can matter most. In one study, tracking every cell in growing hiPSC colonies showed that where and how groups of cells moved predicted how the colony would later organise itself, pinpointing the future organising centres before the tissue-scale pattern was visible (31). This turns a cell's spatial context from a nuisance into a useful predictive signal.

Some studies boil movement down to a few numbers that are easy to measure at scale. Measuring how fast progenitor cells moved, whether they survived, and how they divided across long time-lapses fed models that predict division outcomes over several generations (32). Movement can steer fate as well as forecast it: making cells physically tumble for minutes inside a soft, sliding gel pushed stem cells toward stronger differentiation, acting through forces on the nucleus that change how the DNA is packaged, a concrete case where movement itself biases fate (33).

Models that forecast where cells will go can also cut down on imaging. One such model, a Social-GAN (a network that predicts each cell's next positions from its own recent track and those of its neighbours), was used to extend short tracks into longer predicted ones (34). Tested on prostate cancer cells with and without chemotherapy, and on immune cells attacking tumour cells in an organ-on-a-chip, the forecast tracks matched the real ones over a limited window: measures such as speed, straightness and time spent interacting were statistically indistinguishable, so the forecasts kept the drug- and condition-specific movement differences of the real data. Pushed further into the future, this agreement breaks down.

Taken together, these works point to a simple order: track first, predict second. Reliable fate prediction from movement rests on (i) dependable cell outlining and tracking, (ii) track summaries that capture how persistently and how fast a cell moves, how much it turns, and what its neighbours are doing, and (iii) careful separation of features that merely correlate with fate from those that cause it.

Stepping back across these four signal sources, a few generalisations hold. The earliest and steadiest predictions come less from any single modality than from temporal information: time-resolved features anticipate division, death and lineage commitment hours to generations ahead of molecular markers (2, 8, 17, 27), whereas matched single-frame models tend to recover the cell's current state instead of its future one. Which signal dominates is task-dependent: shape for differentiation, intensity for cell-cycle phase, neighbour density for competition, displacement for lineage choice, which is one reason no single architecture wins outright. Two limitations run through all four sections, and we return to them in detail below: most reported prediction is early classification of an already-committed state, and ground truth is almost always endpoint staining read backwards, so prospective and causal tests (33) remain rare.

Architectural Trade-offs: Matching Models to Biology

The studies above use several kinds of neural network, and it helps to know what each is built to do. A neural network learns to turn an input (here, cell images or measurements) into an output (here, a predicted fate) by adjusting many internal weights on training examples. The families differ in the kind of pattern they capture well: some read the layout of a single image, some read how things change over time, and some compress images into a compact numerical summary without being told the answer in advance. The variety in the field reflects real biological variety in where fate information sits (Fig. 5). No single model is best; the useful question is which one matches the biology of a given problem.

When Convolutional Networks Excel. Convolutional neural networks (CNNs) are built to recognise visual patterns in a single image, wherever those patterns happen to sit. They work best when the fate can be read from how a cell looks in one frame, without needing its history. Sorting cells by cell-cycle phase is the clearest example (16, 22): chromatin texture, nuclear size and cytoplasmic density look visibly different in G1, S and G2/M, and a CNN can read these from a single snapshot. Two properties make CNNs a natural fit for microscopy. First, they detect a feature such as condensed chromatin no matter where it sits in the field of view. Second, they build an image up in layers, from simple edges and textures to recognisable structures like a rounded, dividing cell, close to how a biologist reads the same picture.

CNNs struggle, though, when fate depends on how a cell changes over time instead of how it looks at one instant. Two cells can appear identical in a single frame yet go on to different fates set by their earlier behaviour. This limit is what motivates the next two model families, which handle time directly.

When Recurrent Networks and Transformers Outperform. Recurrent networks (RNNs) and transformers earn their keep when a cell's history resolves what a single image cannot. Blood-progenitor commitment, one of the first fate problems solved with a recurrent (LSTM) network, shows this clearly (8): cells look alike at the moment of decision, yet their past, how fast they moved, how often they divided, whether their shape wavered or held steady, predicts which lineage they choose. An RNN reads frames one after another, carrying a running memory that accumulates what it has seen and learning which past events matter. A transformer does something similar but can look back at all timepoints at once and weigh each by how relevant it is, in principle without being told where to look.

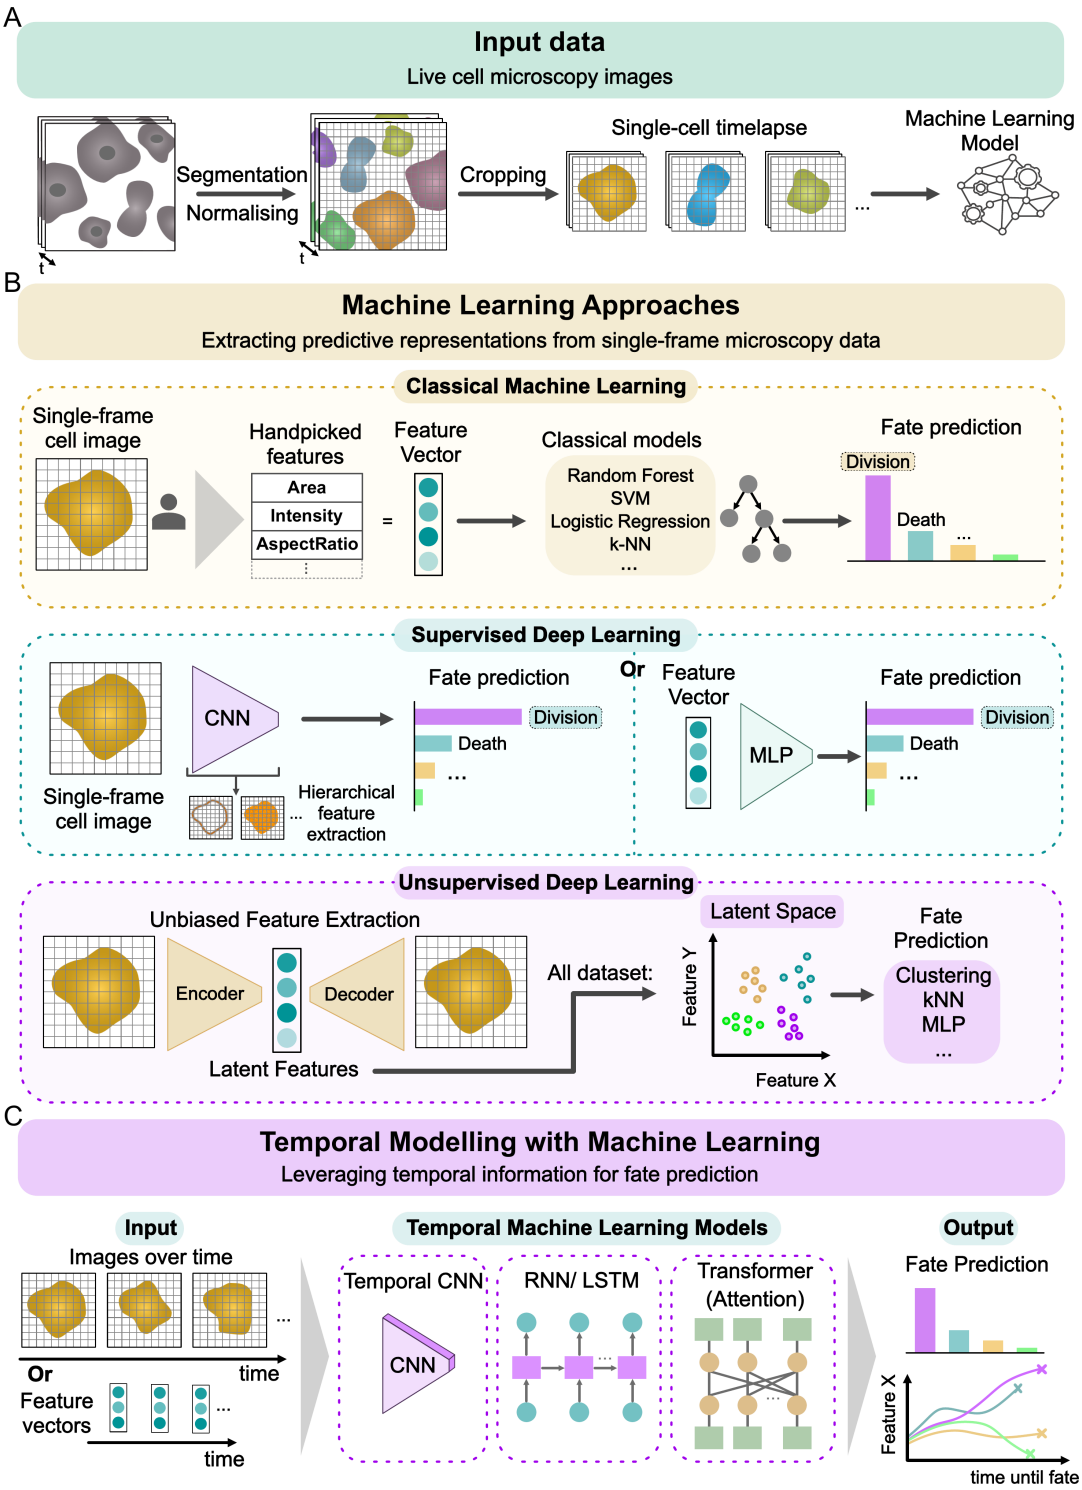

Fig. 5. Machine-learning frameworks for cell fate prediction from microscopy. A. Data pipeline: time-lapse images are outlined into single cells, normalised, and cropped into per-cell image stacks for the models. B. Three ways to analyse single frames: classical machine learning on hand-chosen measurements (methods such as Random Forest, SVM, Logistic Regression, k-NN); supervised deep learning that learns features itself (CNNs or fully connected networks); and unsupervised deep learning (autoencoders) that compresses each image into a compact numerical summary for grouping or later classification. C. To use time, these same approaches are extended to sequences of images or measurements, with networks built to read ordered data (temporal CNNs, recurrent networks/LSTMs, or attention-based transformers), so fate can be predicted from each cell's history.

The cost is data and compute. A CNN needs only a single frame per cell, whereas following cells over time means collecting and labelling image sequences that span hours to days, a much larger burden. Transformers, which have more internal parameters to fit, usually need still more training data to work well and are slower to train. For high-throughput screening or labs with limited resources, a simpler CNN therefore stays attractive even if its best-case accuracy is lower.

Variational Autoencoders: The Unsupervised Strategy. Autoencoders, including VAEs and their variants (β-VAE, τ-VAE, adversarial autoencoders), have a different strength: they learn from unlabelled images. When a cell's fate is unknown until days after imaging, or when labelling enough examples is too costly, these models still find useful structure on their own. Each image is compressed into a short list of numbers, and that compact summary itself becomes a meaningful description of the cell's state (23). Cells then sit as points in a continuous space where nearby points look alike and a path through the space traces how a cell changes over time, instead of being forced into rigid boxes like G1, S or G2/M. This keeps more information than fixed labels while staying easier to interpret than the raw pixels (2, 23).

This flexibility has costs. The summary numbers are not automatically meaningful; each one may mix several unrelated properties unless the model is designed to prevent it. β-VAEs push each number to track one independent property, which helps, though it still does not guarantee that a number matches something a biologist recognises. Temporal VAEs (τ-VAE) keep successive frames close, buying smoother, more stable trajectories at the price of missing abrupt fate switches. Adversarial autoencoders shape the summary space in a chosen way but need careful tuning, or they collapse to producing only a few stereotyped outputs.

In cost, VAEs sit between CNNs and transformers. Turning a new image into its summary is fast, a single pass through the network, though training the model in the first place is more involved. When the fates of interest are not yet defined, VAEs are invaluable for reducing a flood of image measurements to a handful of interpretable numbers. When fates are well defined and plenty of labelled examples exist, a supervised CNN or RNN usually does better.

Toward Hybrid Architectures. Beyond the per-model costs above, methods that treat cells as 3D shapes need heavy computation and specialised 3D imaging, and in high-throughput screening the choice of model is decided as much by inference speed and available hardware as by accuracy.

Future models will likely combine the strengths of each: a CNN to read each image, an RNN or transformer to follow cells over time, and an autoencoder to learn from unlabelled data first. Such hybrids can learn general features from large unlabelled time-lapse collections, then be fine-tuned on the smaller set of labelled examples that exist for a specific fate. This helps with a real bottleneck: labelled fate outcomes are scarce, while unlabelled time-lapse data are abundant. Tracking can also generate labels: follow each cell forward to its fate, then carry that fate back onto the earlier, unlabelled frames, as shown in developing embryos (35). Progress will come less from novel algorithms than from matching a model's built-in assumptions to how fate is actually decided in cells.

Biological Meaning of Morphological Features

A consistent challenge in applying deep learning to biological imaging is that many models are hard to understand. They often reach high accuracy, yet it stays unclear which image features they are using and whether those features mean anything biologically.

The imaging itself can change what cells do. Too much light, for instance, is phototoxic and can disturb a cell's functions, its organelles, and its development. A model may then learn to recognise damage caused by imaging instead of real biology, which is why tools that detect phototoxicity are needed (36).

Avoiding this means checking that the training images reflect real biology instead of imaging damage. PhotoFiTT (37) helps here, giving a label-free, quantitative way to measure phototoxicity in live-cell experiments. By tracking how cells divide, change size and stay active, it lets researchers choose light levels that leave cells behaving normally. Such checks matter for interpretation, because they lower the chance that a model's predictive features come from experimental artefacts instead of biology.

Interpretable models can also reveal which features carry the signal. A model of label-free live-cell images pinpointed the shape properties that mark highly metastatic melanoma, turning a black-box prediction into a specific, testable cell feature (38).

Adding label-free imaging modes such as phase contrast has been proposed to enrich the information available (22), yet ways to interpret these models still need work. Many continue to rely on hidden correlations that are hard to check against biology. Progress in interpretation and feature discovery will take more than fancier models; it will also need careful experimental design, transparent representations, and systematic testing of the features that drive predictions.

What Is Established, and What Is Not

Where the evidence is strong and where it is not deserves a plain statement. Three distinctions matter. First, much of what is called fate *prediction* is in practice early *classification* of a state the cell has already entered; genuine forward prediction, of an outcome not yet determined at the time of imaging, is rarer and harder, and the two are frequently conflated. Second, reconstructing cell-cycle position from images, now achievable from intensity, phase or raw morphology (22, 24–26), shows that state is encoded optically, but inferring a current state is not the same as predicting a future fate. Third, almost all reported performance is validated retrospectively against endpoint markers; prospective validation is uncommon, and causal tests, showing that altering a predicted trajectory changes the outcome, amount to essentially a single example (33). None of this diminishes the progress, but it sets the bar for what would make these tools trustworthy: forward prediction ahead of late classification, prospective validation ahead of retrospective, and features that can be tied to mechanism.

Future Perspectives

Microscopy-based cell fate prediction has so far focused mostly on the cell cycle and cell death, which carry strong morphological signatures and are experimentally accessible. This

established proof of principle, but it leaves other areas of cell biology unexplored. Processes such as migration, differentiation, senescence and stress responses also involve coordinated, time-dependent changes in cell and nuclear morphology, yet stay underrepresented as prediction targets.

Cell migration is an especially promising area, since features such as directional persistence, speed fluctuations, protrusive dynamics and polarisation evolve continuously and tie closely to functional outcomes. These morphodynamic signatures may predict transitions between stationary and migratory states, or between migration modes, before canonical markers appear. Defining fate here stays challenging, though, and calls for long-term, functionally grounded experiments.

Cellular senescence is another important but underexplored fate. Senescent cells show characteristic morphological changes, including nuclear enlargement and altered chromatin organisation, and deep learning can already classify senescence from nuclear morphology (14). Whether such features can flag pre-senescent states before established markers appear stays an open, largely untested question, and telling senescence apart from other non-proliferative states is hard.

Predicting differentiation beyond binary outcomes likewise stays challenging, since early morphological differences between lineages tend to be subtle. High-dimensional, long-term morphodynamic analysis, perhaps with spatial context, may help resolve these trajectories. Stress responses, for their part, induce distinct temporal adaptations, and modelling these dynamics could predict whether cells recover, arrest or die, which would be valuable for toxicology and disease models.

A key direction is linking a cell's appearance to its molecular state. Single-cell methods that read many molecular layers at once, and methods that infer the direction of change from the balance of new and mature RNA, already reconstruct molecular paths toward fate (39, 40), and fate is increasingly predicted straight from gene-expression state, whether by combining several deep models (41), by watching for early-warning signs that a cell is nearing a decision point (42), or by models that fuse different molecular layers (43). Imaging-based prediction works best alongside these molecular methods, not against them: pairing images with molecular data could connect visible features to the mechanisms behind them, helping explain why a given appearance predicts a given fate and suggesting experiments to test it.

Beyond the biology, the field needs shared infrastructure so that results can be compared across labs. Three steps would help. First, standard benchmarks that test a model on cells it has never seen, and predict their fate before it is known, built on existing resources such as the JUMP Cell Painting collection (7) and large 3D single-cell datasets. Second, evaluation that goes beyond raw accuracy to report how well a model's confidence is calibrated and how far ahead of the molecular markers it can call a fate. Third, open repositories of trained models and annotated time-lapse data with common reporting standards. Models that teach themselves reusable features from images with no manual labels (44), along with large pretrained 'foundation' models, may cut the data needed further by letting a general model be fine-tuned on a small labelled set.

Two major challenges remain: interpretability and validation. Current models regularly rest on correlations that are not mechanistically understood, and predictions are usually validated retrospectively instead of prospectively. Moving toward causal insight will mean integrating morphology with molecular data and testing whether predicted fates can be experimentally altered. Together, these advances will push fate prediction past simple endpoints toward a fuller, mechanistic understanding of cell behaviour.

Conclusions

To conclude, the studies reviewed here establish that morphology, and especially its temporal evolution, reflects underlying biophysical and biochemical processes that shape cell-cycle progression, differentiation, survival, and competition. In some settings shape is more than a passive readout of molecular regulation: it can feed back on fate, as when physical forces on the nucleus change how DNA is packaged and, with it, what the cell becomes. Static measurements give only partial insight, whereas temporal trajectories reveal early, non-invasive signatures that anticipate division, death, differentiation, and competitive outcomes well before conventional molecular markers emerge. As imaging technologies and computational models co-evolve, backed by solid open-data infrastructure and standardised model repositories, morphology-driven fate prediction will reach past the cell cycle to migration, senescence, stress responses, and fine-grained differentiation decisions. Key challenges remain despite this progress. The balance between predictive performance and biological interpretability, the integration of multimodal data, and the need to capture both gradual and abrupt state transitions continue to define the limits of current approaches. Working through these challenges will be essential to understand the biological mechanisms behind the predictions. Bringing morphodynamic, molecular, and computational perspectives together offers a path toward a more predictive and mechanistic understanding of cell behaviour, bearing implications for developmental biology, disease modelling, and therapeutic intervention. With further validation and standardisation, morphological fate prediction could become a routine quantitative assay, letting biologists ask when, why, and how cells commit to their futures.

ABOUT THIS MANUSCRIPT

This manuscript was prepared using Rχiv-Maker v1.22.2 (45).

AUTHOR CONTRIBUTIONS

R.C. and M.D.R. led the manuscript structure and writing. R.C., M.D.R., I.C., J.G., G.J. and R.H. contributed to manuscript drafting, literature curation and figure design. G.J. and R.H. supervised the project. All authors revised and approved the final manuscript.

FUNDING

This work was supported by FCT - Fundação para a Ciência e a Tecnologia, I.P., through the MOSTMICRO-ITQB R&D Unit (DOI 10.54499/UID/04612/2025, UID/PRR/4612/2025) and the LS4FUTURE Associated Laboratory (DOI 10.54499/LA/P/0087/2020).

R.H. acknowledges funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (SelfDriv-

ing4DSR, grant agreement No. 101001332). R.H. is further supported by the European Union through Horizon Europe (RT-SuperES, grant agreement No. 101099654); by a European Molecular Biology Organization (EMBO) Installation Grant (EMBO-2020-IG-4734); by a Chan Zuckerberg Initiative Visual Proteomics Imaging award (vpi-0000000044, https://doi.org/10.37921/743590vtudfp); by a joint Wellcome, Chan Zuckerberg Initiative, and Kavli Foundation Essential Open Source Software for Science Cycle 6 award (Wellcome 313383/Z/24/Z; CZI EOSS6-0000000260); and by the "la Caixa" Foundation (CaixaResearch Health 2025, VirusAwareScopes, HR25-00453). R.C. is supported by a "la Caixa" doctoral fellowship under HR25-00453. M.D.R. acknowledges the Center of Excellence IMMENS funded by the Research Council of Finland (374180).

G.J. acknowledges funding from the Research Council of Finland (338537, 371287, and 374180), the Sigrid Jusélius Foundation, the Cancer Society of Finland (Syöpäjärjestöt), and the Solutions for Health strategic funding for Åbo Akademi University, and from the InFLAMES Flagship Programme of the Research Council of Finland (decision numbers 337530, 337531, and 357910). G.J. is supported by the Finnish Cancer Institute (K. Albin Johansson Professorship). I.C. and J.G. acknowledge support from the Knut and Alice Wallenberg Foundation through the Data-Driven Life Science (DDLS) programme (grant 31003604) and project grant 31005835.

Funded by the European Union. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the granting authority. Neither the European Union nor the granting authority can be held responsible for them. The Chan Zuckerberg Initiative awards are made through the Chan Zuckerberg Initiative DAF, an advised fund of Silicon Valley Community Foundation.

COMPETING FINANCIAL INTERESTS

The authors declare no competing interests.

EXTENDED AUTHOR INFORMATION

- **Rita Carlota:**
  0000-0003-3647-1114
- **Mario Del Rosario:**
  0000-0002-0430-1463; ★ mariodelr.bsky.social
- **Inês Cunha:**
  0000-0002-1327-9018; ★ inesmcunha.bsky.social
- **Juliette Griffié:**
  0000-0002-0141-4550; ★ juliettegriffie.bsky.social
- **Guillaume Jacquemet:**
  0000-0002-9286-920X; ★ guijacquemet.bsky.social
- **Ricardo Henriques:**
  0000-0002-2043-5234; ★ henriqueslab.bsky.social

Bibliography

1. Jeremy Copperman, Sean M. Gross, Young Hwan Chang, Laura M. Heiser, and Daniel M. Zuckerman. Morphodynamical cell state description via live-cell imaging trajectory embedding. *Communications Biology*, 6:484, 2023. doi:10.1038/s42003-023-04837-8.
2. Christopher J. Soelistyo, Giulia Vallardi, Guillaume Charras, and Alan R. Lowe. Learning biophysical determinants of cell fate with deep neural networks. *Nature Machine Intelligence*, 4:636–644, 2022. doi:10.1038/s42256-022-00503-6.
3. Trang Le, William D. Leineweber, Matheus P. Viana, Andrew Cesnik, Jan N. Hansen, Wei Ouyang, Susanne M. Rafelski, and Emma Lundberg. Cell shapes decode molecular phenotypes in image-based spatial proteomics. *Cell Systems*, 2026. doi:10.1016/j.cels.2026.101589.
4. Anna Ligasová, Ivo Frydrych, and Karel Koberna. Basic methods of cell cycle analysis. *International Journal of Molecular Sciences*, 24, 2023. doi:10.3390/ijms24043674.
5. G. Follain, M. Dibus, O. Joshi, and G. Jacquemet. Time, the final frontier. *Molecular Oncology*, 19(8):2157–2162, 2025. doi:10.1002/1878-0261.70025.
6. Mark-Anthony Bray, Shantanu Singh, Han Han, Chadwick T. Davis, Blake Borgeson, Cathy Hartland, Maria Kost-Alimova, Sigrun M. Gustafsdottir, Christopher C. Gibson, and Anne E. Carpenter. Cell painting, a high-content image-based assay for morphological profiling using multiplexed fluorescent dyes. *Nature Protocols*, 11(9):1757–1774, 2016. doi:10.1038/nprot.2016.105.
7. Srinivas Niranj Chandrasekaran, Beth A. Cimini, Amy Goodale, Lisa Miller, Maria Kost-Alimova, Nasim Jamali, John G. Doench, et al. Three million images and morphological profiles of cells treated with matched chemical and genetic perturbations. *Nature Methods*, 21(6):1114–1121, 2024. doi:10.1038/s41592-024-02241-6.
8. F. Buggenthin, F. Buettner, P. S. Hoppe, M. Endele, M. Kroiss, M. Strasser, M. Schwarzfischer, T. Schroeder, C. Marr, and F. J. Theis. Prospective identification of hematopoietic lineage choice by deep learning. *Nature Methods*, 14:403–406, 2017. doi:10.1038/nmeth.4182.
9. Eric M. Christiansen, Samuel J. Yang, D. Michael Ando, Ashkan Javaherian, Gaia Skibinski, Scott Lipnick, Elliot Mount, Alison O'Neil, Kevan Shah, Alicia K. Lee, Piyush Goyal, William Fedus, Ryan Poplin, Andre Esteva, Marc Berndl, Lee L. Rubin, Philip Nelson, and Steven Finkbeiner. In silico labeling: Predicting fluorescent labels in unlabeled images. *Cell*, 173(3):792–803.e19, 2018. doi:10.1016/j.cell.2018.03.040.
10. Chawin Ounkomol, Sharmishtaa Seshamani, Mary M. Maleckar, Forrest Collman, and Gregory R. Johnson. Label-free prediction of three-dimensional fluorescence images from transmitted-light microscopy. *Nature Methods*, 15:917–920, 2018. doi:10.1038/s41592-018-0111-2.
11. Erick Moen, Dylan Bannon, Takamasa Kudo, William Graf, Markus Covert, and David Van Valen. Deep learning for cellular image analysis. *Nature Methods*, 16:1233–1246, 2019. doi:10.1038/s41592-019-0403-1.
12. M. Tegtmeyer, J. Arora, S. Asgari, B. A. Cimini, A. Nadig, E. Peirent, D. Liyanage, G. P. Way, et al. High-dimensional phenotyping to define the genetic basis of cellular morphology. *Nature Communications*, 15:347, 2024. doi:10.1038/s41467-023-44045-w.
13. G. P. Way, T. Natoli, A. Adeboye, L. Litichevskiy, A. Yang, X. Lu, J. C. Caicedo, and B. A. Cimini. Morphology and gene expression profiling provide complementary information for mapping cell state. *Cell Systems*, 13(11):911–923.e9, 2022. doi:10.1016/j.cels.2022.10.001.
14. Ina Heckenbach, Garik V. Mkrtchyan, Mikolaj B. Ezra, Dominika Bakula, Jesse S. Madsen, Malte H. Nielsen, Daniela Oró, Bethany Osborne, Anthony J. Covarrubias, Maria Laura Idda, Sofie Lautrup, Henok Kassahun, and Morten Scheibye-Knudsen. Nuclear morphology is a deep learning biomarker of cellular senescence. *Nature Aging*, 2:742–755, 2022. doi:10.1038/s43587-022-00263-3.
15. Laura Wiggins, Alice Lord, Killian L. Murphy, Stuart E. Lacy, Peter J. O'Toole, William J. Brackenbury, and Julie Wilson. The CellPhe toolkit for cell phenotyping using time-lapse imaging and pattern recognition. *Nature Communications*, 14:1864, 2023. doi:10.1038/s41467-023-37447-3.
16. Yukiko Nagao, Mika Sakamoto, Takumi Chinen, Yasushi Okada, and Daisuke Takao. Robust classification of cell cycle phase and biological feature extraction by image-based deep learning. *Molecular Biology of the Cell*, 31:1346–1354, 2020. doi:10.1091/mbc.E20-03-0187.
17. Mari Tolonen, Ziwei Xu, Ozgur Beker, Varun Kapoor, Bianca Dumitrascu, and Jakub Sedzinski. Single-cell morphodynamics predict cell fate decisions during mucociliary epithelial differentiation. *bioRxiv*, 2025. doi:10.1101/2025.09.17.676780. Preprint.
18. Maxwell Mai, Shuai Luo, Samantha Fasciano, Timilehin Esther Oluwole, Justin Ortiz, Yulei Pang, and Shue Wang. Morphology-based deep learning approach for predicting adipogenic and osteogenic differentiation of human mesenchymal stem cells (hmscs). *Frontiers in Cell and Developmental Biology*, 11:1329840, 2023. doi:10.3389/fcell.2023.1329840.
19. Matt De Vries, Lucas G. Dent, Nathan Curry, Adam Tyson, Chris Dunsby, and Chris Bakal. Geometric deep learning and multiple-instance learning for 3D cell-shape profiling. *Cell Systems*, 16:101229, 2025. doi:10.1016/j.cels.2025.101229.
20. Rory M. Donovan-Maiye, Jackson M. Brown, Caleb K. Chan, Liya Ding, Calysta Yan, Nathalie Gaudreault, Julie A. Theriot, Mary M. Maleckar, Theo A. Knijnenburg, and Gregory R. Johnson. A deep generative model of 3d single-cell organization. *PLOS Computational Biology*, 18(1):e1009155, 2022. doi:10.1371/journal.pcbi.1009155.
21. Inês Cunha, Luca Panconi, Sebastian Bauer, Maxime Gestin, Emma Latron, Erik Sahai, Alix Le Marois, and Juliette Griffié. AI4CellFate: Interpretable early cell fate prediction with generative AI. *bioRxiv*, 2025. doi:10.1101/2025.05.12.653464. Preprint.
22. Yuchen R. He, Shenghua He, Mikhail E. Kandel, Young Jae Lee, Chenfei Hu, Nahil Sobh, Mark A. Anastasio, and Gabriel Popescu. Cell cycle stage classification using phase imaging with computational specificity. *ACS Photonics*, 9:1264–1273, 2022. doi:10.1021/acsphotonics.1c01779.
23. Thomas Bonte, Oriane Pourcelot, Adham Safieddine, Floric Slimani, Florian Mueller, Dominique Weil, Edouard Bertrand, and Thomas Walter. A deep learning approach for time-consistent cell cycle phase prediction from microscopy data. *bioRxiv*, 2025. doi:10.1101/2025.05.16.654306. Preprint.
24. Philipp Eulenberg, Niklas Köhler, Thomas Blasi, Andrew Filby, Anne E. Carpenter, Paul Rees, Fabian J. Theis, and F. Alexander Wolf. Reconstructing cell cycle and disease progression using deep learning. *Nature Communications*, 8:463, 2017. doi:10.1038/s41467-017-00623-3.
25. Luca Rappez, Alexander Rakhlin, Angelos Rigopoulos, Sergey Nikolenko, and Theodore Alexandrov. Deepcycle reconstructs a cyclic cell cycle trajectory from unsegmented cell images using convolutional neural networks. *Molecular Systems Biology*, 16(10):e9474, 2020. doi:10.15252/msb.20209474.
26. Diana Mahdessian, Anthony J. Cesnik, Christian Gnann, Frida Danielsson, Lovisa Stenström, et al. Spatiotemporal dissection of the cell cycle with single-cell proteogenomics. *Nature*, 590(7847):649–654, 2021. doi:10.1038/s41586-021-03232-9.
27. Alexis J. Chambost, Nabila Berabez, Olivier Cochet-Escartin, François Ducray, Mathieu Gabut, Caroline Isaac, Sylvie Martel, Ahmed Idbaih, David Rousseau, David Meyronet, and Sylvain Monnier. Machine learning-based detection of label-free cancer stem-like cell fate. *Scientific Reports*, 12:19066, 2022. doi:10.1038/s41598-022-21822-z.
28. Sarah He, Muhammed Sillah, Aidan R. Cole, Apoorva Uboveja, Katherine M. Aird, Yu-Chih Chen, and Yi-Nan Gong. D-MAINS: A deep-learning model for the label-free detection of mitosis, apoptosis, interphase, necrosis, and senescence in cancer cells. *Cells*, 13(12):1004, 2024. doi:10.3390/cells13121004.
29. Edward Ren, Sungmin Kim, Saad Mohamad, Samuel F. Huguet, Yulin Shi, Andrew R. Cohen, Eugenia Piddini, and Rafael Carazo Salas. Deep learning-enhanced morphological profiling predicts cell fate dynamics in real-time in hpscs. *bioRxiv*, 2021. doi:10.1101/2021.07.31.454574. Preprint.
30. Michael J. Shannon, Shira E. Eisman, Alan R. Lowe, Tyler F. W. Sloan, and Emily M. Mace. cellPLATO – an unsupervised method for identifying cell behaviour in heterogeneous cell trajectory data. *Journal of Cell Science*, 137:jcs261887, 2024. doi:10.1242/jcs.261887.
31. David A. Joy, Ashley R. G. Libby, and Todd C. McDevitt. Deep neural net tracking of human pluripotent stem cells reveals intrinsic behaviors directing morphogenesis. *Stem Cell Reports*, 16:1317–1330, 2021. doi:10.1016/j.stemcr.2021.04.008.
32. Vanessa M. Scanlon, Ana M. Teixeira, Taru Tyagi, Siyan Zou, Peng Zhang, Christopher A. J. Booth, Patricia Cazenave, Seema Barthwal, Diana Odru, Andrea Pellagatti, Jacqueline Boultwood, Tudor A. Fulga, and Irene Roberts. Multiparameter analysis of timelapse imaging reveals kinetics of megakaryocytic erythroid progenitor clonal expansion and differentiation. *Scientific Reports*, 12:16159, 2022. doi:10.1038/s41598-022-19013-x.
33. Manish Ayushman, Georgios Mikos, Xin Tong, David Mun, Shaochen Dai, Edward Im, Anusorn Kannan, James Zou, and Fan Yang. Cell tumbling enhances stem cell differentiation in hydrogels via nuclear mechanotransduction. *Nature Materials*, 24:312–322, 2025. doi:10.1038/s41563-024-02038-0.
34. M. C. Comes, J. Filippi, A. Mencattini, F. Corsi, P. Casti, A. De Ninno, et al. Accelerating the experimental responses on cell behaviors: a long-term prediction of cell trajectories using social

generative adversarial network. *Scientific Reports*, 10:15635, 2020. doi:10.1038/s41598-020-72605-3.

35. Yinan Wan, Jakob El Kholtei, Ignatius Jenie, Mariona Colomer-Rosell, Jialin Liu, Qinghua Zhang, Joaquin Navajas Acedo, Lucia Y. Du, Mireia Codina-Tobias, Mengfan Wang, Wei Zheng, Edward Lin, Tzy-Harn Chuang, Oded Mayseless, Ahilya Sawh, Susan E. Mango, Guoqiang Yu, Bogdan Bintu, and Alexander F. Schier. Whole-embryo spatial transcriptomics at subcellular resolution from gastrulation to organogenesis. *Science*, 2026. doi:10.1126/science.adt3439.
36. David Richmond, Anna Payne-Tobin Jost, Talley Lambert, Jennifer Waters, and Hunter Elliott. Deadnet: Identifying phototoxicity from label-free microscopy images of cells using deep convnets, 2017.
37. Mario Del Rosario, Estibaliz Gómez-de Mariscal, Leonor Morgado, Raquel Portela, Guillaume Jacquemet, Pedro M. Pereira, and Ricardo Henriques. PhotoFiTT: A quantitative framework for assessing phototoxicity in live-cell microscopy experiments. *bioRxiv*, 2024. doi:10.1101/2024.07.16.603046. Preprint.
38. Assaf Zaritsky, Andrew R. Jamieson, Erik S. Welf, Andres Nevarez, Justin Cillay, Ugur Eskiocak, Brandi L. Cantarel, and Gaudenz Danuser. Interpretable deep learning uncovers cellular properties in label-free live cell images that are predictive of highly metastatic melanoma. *Cell Systems*, 12(7):733–747, 2021. doi:10.1016/j.cels.2021.05.003.
39. Laleh Haghverdi and Leif S. Ludwig. Single-cell multi-omics and lineage tracing to dissect cell fate decision-making. *Stem Cell Reports*, 18:13–25, 2023. doi:10.1016/j.stemcr.2022.12.003.
40. Alexander Aivazidis, Fani Memi, Vitalii Kleshchevnikov, Sezgin Er, Brian Clarke, Oliver Stegle, and Omer Ali Bayraktar. Cell2fate infers RNA velocity modules to improve cell fate prediction. *Nature Methods*, 2025. doi:10.1038/s41592-025-02608-3.
41. Jiantao Shen, Nan Chen, Bowen Niu, Jiayuan Zhong, and Rui Liu. FatePredictor: Cell fate decision-making prediction with an ensemble deep learning model. *The Innovation*, 6(11): 101010, 2025. doi:10.1016/j.xinn.2025.101010.
42. Mehrshad Sadria and Thomas M. Bury. FateNet: an integration of dynamical systems and deep learning for cell fate prediction. *Bioinformatics*, 40(9):btae525, 2024. doi:10.1093/bioinformatics/btae525.
43. Wenzhuo Tang, Hongzhi Yin, Renming Mao, and Karin Verspoor. Single-cell multimodal prediction via transformers. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 2436–2446, 2023. doi:10.1145/3583780.3615061.
44. Hirofumi Kobayashi, Keith C. Cheveralls, Manuel D. Leonetti, and Loic A. Royer. Self-supervised deep learning encodes high-resolution features of protein subcellular localization. *Nature Methods*, 19(8):995–1003, 2022. doi:10.1038/s41592-022-01541-z.
45. Bruno M. Saraiva, Rita Carlota, António D. Brito, Iván Hidalgo-Cenalmor, Guillaume Jacquemet, and Ricardo Henriques. Rxiv-maker: an automated template engine for streamlined scientific publications, 2025. doi:10.48550/arXiv.2508.00836.