

Graph-based machine learning reveals rules of spatiotemporal cell interactions in tissues

Takaki Yamamoto,^{1,*} Katie Cockburn,² Valentina Greco,^{2,3} and Kyogo Kawaguchi^{1,4,5,†}

¹*Nonequilibrium Physics of Living Matter RIKEN Hakubi Research Team,
RIKEN Center for Biosystems Dynamics Research, 2-2-3 Minatojima-minamimachi, Chuo-ku, Kobe 650-0047, Japan*

²*Department of Genetics, Yale School of Medicine, New Haven, CT 06510, USA*

³*Departments of Cell Biology and Dermatology, Yale Stem Cell Center,
Yale Cancer Center, Yale School of Medicine, New Haven, CT 06510, USA*

⁴*RIKEN Cluster for Pioneering Research, 2-2-3 Minatojima-minamimachi, Chuo-ku, Kobe 650-0047, Japan*

⁵*Universal Biology Institute, The University of Tokyo, Bunkyo-ku, Tokyo 113-0033, Japan*

(Dated: June 23, 2021)

Extracting the rules of cell-to-cell interactions in tissue dynamics is challenging even if high-resolution live microscopy is accessible. In order to seek and compare the different rules enforced in developing and homeostatic tissues, it will be desirable to have a systematic method that outputs the rules of multi-cellular kinetics simply from live images and cell tracks. Here we demonstrate that graph neural network (GNN)-based models can predict cell fate in the mammalian epidermis when spatiotemporal graph constructed from cell tracks and contacts are given as inputs. By extracting the rules learned by GNN, we find neighbor cell fate inductions and inhibitions consistent with previous findings as well as some that have been previously overlooked. This study demonstrates how GNN-based methods can be useful in inferring stochastic dynamics such as multi-cellular kinetics involving proliferation and loss of agents.

INTRODUCTION

One of the grand challenges in the physical study of biology is to write down the “equation of motion” that governs the time evolution of whole multicellular systems. Recent advances in single-cell profiling experiments and computational methods have refined the molecular view on cell differentiation processes [1] and enabled the prediction of cell lineage trees at a massive scale [2, 3]. Yet, it is still difficult to assess the rules of cell-to-cell interactions, which are known to play key roles in all stages of multicellular dynamics [4]. Although there are promising attempts to probe the cell interactions from high-throughput data [5–11], the relevance of the interactions in the spatiotemporal domain (e.g., their effects on determining the fate of the cells) can so far only be addressed through direct live imaging.

In the homeostasis of the adult mammalian epidermis, where the basal layer cells undergo rapid turnover within a time scale of days, clonal analyses have shown that cell divisions and differentiations are stochastic [12, 13]. It has recently been further identified through live imaging that there exists a correlation in the fates between the neighboring cells, indicating that the basal layer cells are communicating with each other to maintain the stem cell pool [14]. Similar mechanisms have been found in the drosophila intestine [15] as well as in the mammalian intestine [16] and stomach [17]. In general, the stem cell pool in homeostatic tissues must be maintained through some feedback mechanism acting between the cells [18], either directly through cell contacts [14–17] or in more indirect forms such as niche competition [19, 20] in order to prevent overgrowth or depletion.

Inferring the rules of interaction from experiments in tissue dynamics typically requires case-by-case analyses involving

the testing of specific hypotheses based on live imaging and perturbation experiments. Alternatively, it will be desirable to have a computational pipeline where rules behind the dynamics of the tissue can be unbiasedly suggested after feeding in live image data and cell tracks. One way will be to train a model based on machine learning to predict the future dynamics of cells from past data, and later challenge the model by partially hiding data, applying attribution methods, or using symbolic reduction to extract the rules that have been learned by the machine.

Machine learning methods have been heavily used to extract features of cells from live or fixed cell image sequences. Applications include segmentation and tracking [21–23], in silico labeling [24], and direct prediction of cell fates [25], which all use versions of convolutional neural networks (CNN). The downside in using image data directly for cell fate prediction is that images typically contain redundant information making it difficult to focus on the most relevant components to interpret the rules. For instance, a key feature of cells undergoing cell division is their size growth, which means that the machine will likely learn through images that cell size is a good predictor of cell fate. This strong association will be a problem when we are interested in the mechanism upstream of fate decision (i.e., commitment to cell division), since cell size is difficult to mask out from images. On the other hand, important aspects such as cell tracks, lineages, and contacts are not trivially deducible by the machines from the images, making it harder to conduct interpretable rule extractions based on the experiences of individual cells.

An alternative machine learning approach toward automatic rule discoveries is the graph neural network (GNN) scheme [26]. The GNN approach has been utilized not only in forward dynamics predictions of physical systems [27] but has also been successful in inferring the rules of agent-interactions [28] and dynamical properties [29] from time-lapse data obtained by simulations and experiments. The GNN framework is particularly suited for characterizing the

* takaki.yamamoto@riken.jp

† kyogo.kawaguchi@riken.jp

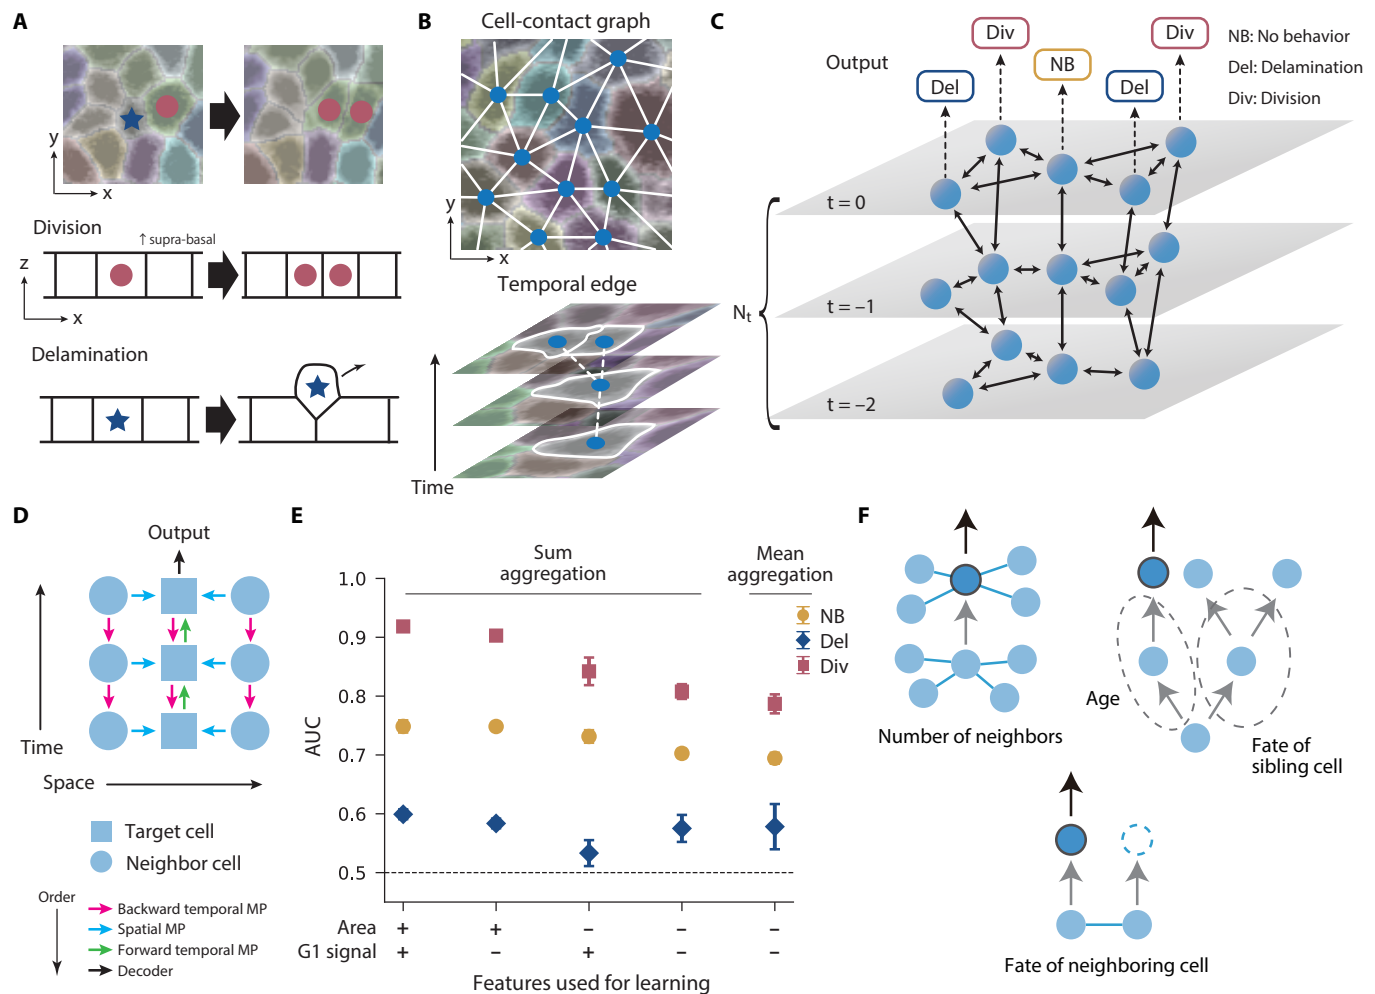


Fig. 1. Bidirectional spatiotemporal GNN model: analysis of the hind paw epidermis data. (A) Schematic of cell division and delamination in the basal layer of the mouse epidermis. (B) Schematic of the cell-contact graph and temporal edges in the basal layer. (C) Schematic of the 2+1 dimensional spatiotemporal graph. A 3-time model is shown as an example. The spheres represent cells, and double-headed arrows represent the neighboring cell connection and the temporal tracks. The GNN models predict the cell fate (NB, Del, Div) conducted in the final layer ($t = 0$). (D) Schematic of the information flow in the bidirectional spatiotemporal GNN model. The information flow in a target cell's subgraph is shown. The target and the neighboring cells are represented by boxes and circles, respectively. Each arrow indicates the direction of the information flow, the order of which is represented by the different colors. (E) The area under the curve score (AUC) of the bidirectional spatiotemporal GNN models with sum and mean aggregation is shown for models under various feature conditions. The AUC for each cell fate label is shown with different markers (NB: circle, Del: diamond, Div: square), which are obtained by averaging over six trained models. Error bar: standard deviation. (F) Graph motifs that are potentially useful in predicting the target cell fate. The correlation of these motifs to the target cell fate can be exploited to make predictions in the bidirectional spatiotemporal GNN model.

dynamics in multicellular systems since it can represent the heterogeneous contact interactions between cells, temporal cell tracks involving occasional cell division and loss, together with the intracellular features such as gene expressions or any other information extractable from image data. Representing track data of a multicellular tissue as spatiotemporal graphs is also a natural way of data compression; the characteristics of the graphs can be compared across tissues and organisms as well as to graphs generated by numerical models and from contexts outside of biology [30]. In this work, we show how the GNN-based approach can predict the fates of cells in the

epidermal stem cell pool, and demonstrate how the rules of cell interactions can be inferred by systematically reducing the components of GNN models and applying the attribution method.

RESULTS

The epidermis is maintained by continuous cell divisions occurring in a pseudo-two-dimensional region called the basal layer. Cells in the basal layer can irreversibly delaminate

toward the suprabasal layer and differentiate (i.e., turn post-mitotic), and eventually shed off (Fig. 1A). We use the dataset of cell tracks previously generated from live images collected from the non-hairy mouse plantar (hind paw) skin [14], which includes the tracks of all the cells within a region of the basal layer. Between the time frames, cells conduct one of the three possible behaviors; divide (Div), delaminate from the basal layer and migrate into the suprabasal layer (Del), or stay (no behavior, NB), typically with the ratio of 1:1:8 within the time interval of the original image acquisition (12 hours). Loss by cell death is negligibly rare in this tissue under homeostasis. The relative motion between the cells is slow, meaning that the main source of the cell motion is displacement due to cell division and cell delamination.

We construct graphs with the basal layer cells as nodes, connected by edges in spatial and temporal directions. The spatial edges represent the cell-to-cell contact obtained from the two-dimensional segmentation in the basal layer, based on the membrane reporter signal (Fig. 1B). The temporal edges are the tracks of the cells, involving forks corresponding to cell divisions and terminal ends representing cell delamination; we do not track the cells after they have left the basal layer. In each node (i.e., cell in a time frame), multiple features extracted from the raw images can be assigned. For the hind paw data, we chose the cross-section area of the cell in the two-dimensional basal layer calculated from the segmentation, and the level of the Fucci-G1 reporter [31] which is the integral of the corresponding fluorescent channel within the segmented cell area. The GNN model takes in the spatiotemporal graphs as well as the features associated with each node, and outputs predictions on the behavior of cells in the last frame (Div, Del, or NB).

We trained a GNN model using training and test data taken from two separate regions of the same mouse hind paw, each involving more than 200 cells tracked over fifteen time frames with a total of approximately 250 Div and Del events occurring in each region. In the training of a N_t -time model, we extracted graphs composed of sequential N_t time frames, where the time frames are indexed from $t = -(N_t - 1)$ to 0 in temporal order (Fig. 1C), and conducted supervised machine learning to predict the fate of the target cell in the last frame $t = 0$.

The model calculates the output from the input by sequentially updating the features in the nodes in the following way (information flow schematically shown in Fig. 1D: bidirectional spatiotemporal GNN model). First, the model performs message passing (MP) on the temporal edges from the future to the past $N_t - 1$ times. Here, a single MP is conducted by concatenating the features of adjacent nodes to be processed through a multi-layer perceptron (MLP), and then passing on that processed feature to the past node. Second, the model performs MP bidirectionally on the spatial edges N_s times (fig. S1A). Third, the model performs MP $N_t - 1$ times on the temporal edges from the past towards the future nodes. In all MP, we use the sum aggregation in order to properly reflect the graph structure [32]. The order of MP is critical in sending the information of the graph structure. In the end of the MP, the node feature in the last frame incorporates the information from its N_t -frame ancestor nodes, N_s -step neighbor

nodes of the ancestor nodes in the cell contact graph, as well as the daughter cells of the ancestor neighbors (Fig. 1C). We fix $N_s = 1$ in the following analyses since the performance of the model were similar for $N_s = 1$ and 2 (fig. S1), and also $N_t = 4$, which will be changed later. Finally, we decode the node feature in the last frame via another MLP and output the softmax score for each cell fate label (Div, Del, and NB).

The GNN model was successfully trained using standard methods employing a loss function (softmax-cross-entropy loss for the three cell behavior labels) and the Adam optimizer (fig. S2). The Area Under the Curve (AUC) score shows that all behaviors are predicted significantly better than a random guess (AUC=0.5) (Fig. 1E). In particular, we found that cell divisions can be predicted with high accuracy compared with Del and NB. We further found that by training the model with reduced features (i.e., without the cell area, G1-phase reporter signal, or both), the score decreased significantly, indicating that these features are utilized in making reliable predictions. Indeed, the cross-section area tends to be larger for dividing cells (figs. S3A and E), consistent with previous observations [14, 33]. Furthermore, the G1-phase signal tends to be lower in dividing cells as expected, while there was no difference in the G1-phase signal between NB and Del cells (fig. S3F).

Intriguingly, we found that even without the node features, the GNN model can predict the fates with significant scores (Fig. 1E). This indicates that the graph structure itself encodes useful information in predicting cell fate. A candidate structure in the graph that can be used in the prediction is the number of edges (i.e., neighboring cells, Fig. 1F). Indeed, the number of neighboring cells is positively correlated with cell division in the next time frame, which is explained by the fact that cells with a larger area tend to be in contact with more cells in the basal layer (fig. S3C). To see whether the number of edges is important in the prediction, we changed the model to take mean aggregation in the calculation step of the spatial edges, which will make the number of contacts invisible in the model. We found that this change in the model only has a minor effect on the AUC score, suggesting that there are other subgraph motifs that are utilized in the prediction.

Other subgraph motifs which can be interpreted in the biological context include the temporal length up to the last cell division point (Fig. 1F). This is essentially the age of the target cell, which can be used to predict fate if the fate choice is temporally non-random [34]. The neighbor cell fates are also reflected as motifs, which should be important according to previous results [14]. Similarly, the sibling cell behavior can be exploited if there are sibling fate correlations such as in the case of asymmetric division [35], although this is less relevant for the current dataset [14, 34].

Next, we considered reducing the GNN model in order to deduce the effect of graph motifs in the predictions more efficiently. By excluding the initial step of propagating the information toward the past frames, the flow of information becomes unidirectional, which is distinct from the previous model in that the information of the temporal branches in the neighboring cells is not transmitted to the target cell. To compensate for this lack of information, we introduced an addi-

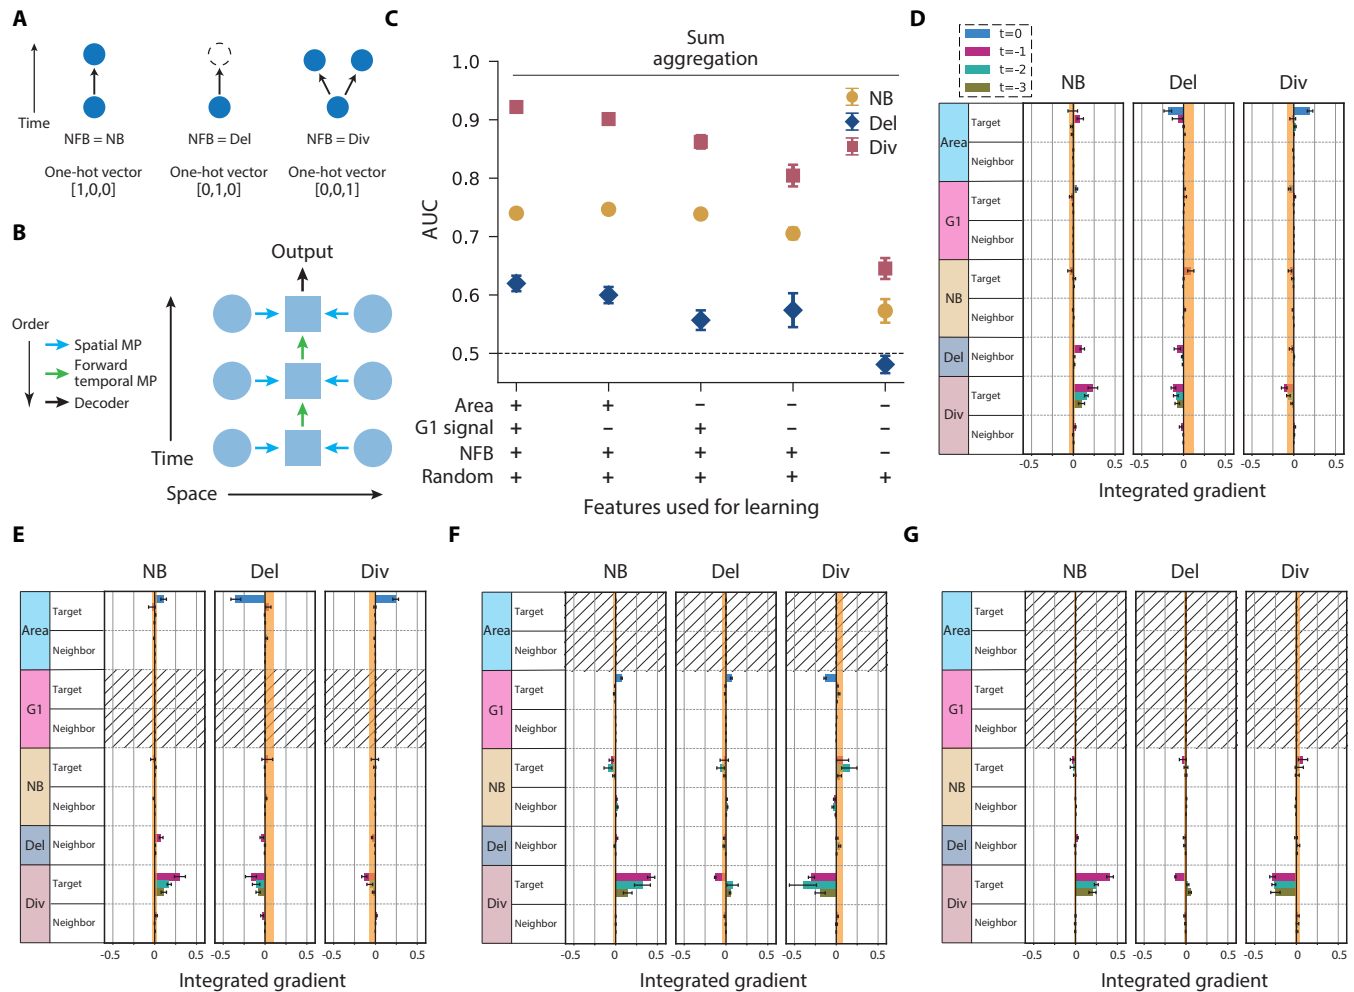


Fig. 2. Unidirectional spatiotemporal GNN model with sum aggregation. (A) The next frame behavior (NFB) feature introduced to represent the local branch structure. (B) Schematic of the information flow in the model of the unidirectional spatiotemporal GNN (see Fig. 1D). (C) The AUC of the unidirectional spatiotemporal GNN is shown for models under various feature sets, obtained by averaging over six trained models. Error bar: standard deviation. (D-G) The attributions of the unidirectional spatiotemporal GNN model are shown for each feature condition: (Area, G1-phase signal, NFB, Random) = (+, +, +, +) (D), (+, -, +, +) (E), (-, +, +, +) (F), and (-, -, +, +) (G). The IG averaged over six trained models is shown for each pooled feature. Error bar: standard error. The upper and lower values of the IGs of the random features are shown as the orange zone.

tional feature to each cell, the next frame behavior (NFB), which takes Div, Del, or NB (Fig. 2A). The reduced GNN model is essentially turning a network structure (i.e., branch and termination representing Div and Del) into a feature of the node, which is useful in addressing the impact of the structure on the predictions. This additional feature was added only to the past cells ($t \leq -1$).

The new GNN model predicts cell fates almost as well as the previous model with and without the area and G1-phase reporter signal features (Fig. 2C and fig. S4), indicating that the key ingredients are captured without the backward-time propagation. By further eliminating the NFB, we confirmed that this feature has an impact on the prediction of all behaviors (Fig. 2C and fig. S5).

So far, we have been treating the GNN as a black box that

outputs predictions of cell fate with varying reliability depending on the inputs. To address the detail of the mechanism of prediction, we next employed the attribution method called the integrated gradient (IG) [36]. In this method, the impact of each input variable on the output is quantified by integrating the change in the output score level upon the gradual shift of each input feature (Figs. 2D-G). The IG also reflects the sign of the impact of each feature; a negative attribution score of feature f toward fate a means that the observation of f will decrease the output score towards predicting a . To account for the noise in IG induced by finite sample size, we further added a random feature to the cells that have no correlation with cell fate and calculated its IG score to identify the baseline for non-zero signal (orange shades, Figs. 2D-G).

First, in Figs. 2D and E, the area of target cells at $t = 0$

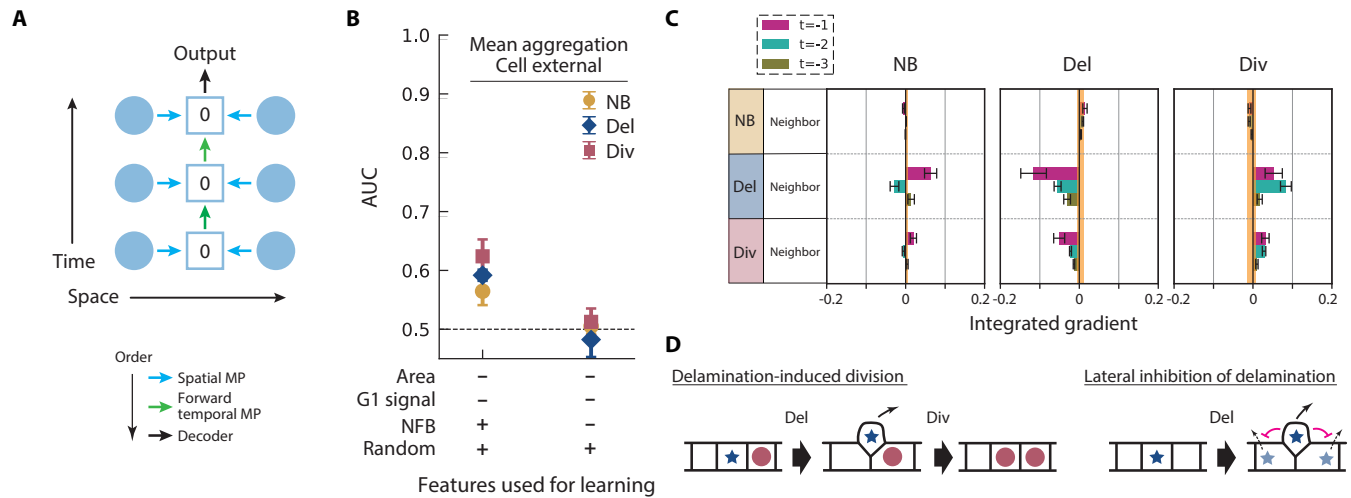


Fig. 3. Cell external model with mean aggregation. (A) Schematic of the information flow in the cell external model (see Fig. 1D). The zero value in the target cell's box represents the null feature vector, which is assigned to the target cell's feature in the cell external model. (B) The AUC of the cell external model with mean aggregation is shown for models with and without future fate, obtained by averaging over six trained models. Error bar: standard deviation. (C) The attribution of the cell external model with mean aggregation is shown for the feature condition: (Area, G1 signal, NFB, Random) = (-, -, +, +). The IG averaged over six trained models is shown for each pooled feature. Error bar: standard error. The upper and lower values of the IG of the random feature are shown as the orange zone. (D) Schematic of a delamination-induced division and lateral inhibition of delamination by neighbor cell delamination.

are found to give positive and negative attribution for the prediction of cell division and delamination, respectively. This indicates that cells with large areas tend to conduct cell division in the next frame, while they are unlikely to delaminate. On the other hand, in Fig. 2F, negative attribution is observed in the G1-phase signal of the target cells at $t = 0$, which is consistent with the expectation that a high G1-phase signal of the target cell should predict no cell division in the next frame. Interestingly, the G1-phase signal has a relatively low score when the cell area feature is present (Fig. 2D); the G1-phase signal is mostly redundant for cell fate prediction. These results highlight how the model is efficiently focusing on the most important variables from the multi-dimensional input.

Without the cell area and G1-phase reporter signal, non-zero attribution scores are found for the prediction of fates by the division of the target cell (Fig. 2G). This is reflecting how the age of the cell can be used in the cell fate prediction; newly born cells tend to undergo a refractory period [34], meaning that the target cell is less likely to divide again if it divided (was born) recently. More generally, the distributions of lifetime can be distinct between cells with different fates (fig. S6), which can be exploited in making fate predictions. To eliminate the cell age information altogether, we set all the features of the target cell sequence to zero (Fig. 3A: cell external model). We also took mean aggregation in conducting the MP on spatial edges to make the number of contacts invisible. This new model can still learn to predict the fates significantly better than random (Fig. 3B). We found a high IG score indicating that Del of a neighboring cell 24 hours before (two frames prior) is useful in predicting Div (Fig. 3C). Consistent with this, when changing the number of input time

frames N_t , the AUC score largely increased between $N_t = 2$ and 3 (fig. S7A). These results match with the previous observation that cell delamination correlates with neighboring cell division 1-2 days later [14].

Moreover, the GNN model predicts that there is an effect of neighboring Del suppressing Del in the next frame (Fig. 3C and D). This indicates that cell delamination is not entirely random, and there may exist a mechanism to suppress two or more neighboring cells to delaminate at the same time. This finding demonstrates that the rule extraction procedure using the GNN model is useful in predicting unexplored mechanisms.

To confirm that the combination of model prediction and the attribution method is indeed pulling out real mechanisms, we conducted numerical simulations of coordinated cell fate models and tested whether artificially implemented rules can be correctly extracted. In the simulated models, cells are represented by points in the 2D space undergoing repulsive interaction with each other. Cell divisions and delaminations are recapitulated by abrupt point duplication (with small noise) and elimination, respectively (Fig. 4A). We encoded the cell fate coordination in the following ways (Fig. 4B). In the delamination-induced division setup, we first randomly picked a cell to delaminate and chose another cell randomly from its six-nearest neighbors to divide two frames after the first cell delaminated. In the division-induced delamination setup, the rules were flipped; a randomly picked cell was set to divide while one of its six nearest neighbors was randomly chosen to delaminate after the division. We also tried another setup where these two rules were equally mixed.

Running the same GNN learning algorithm, we found

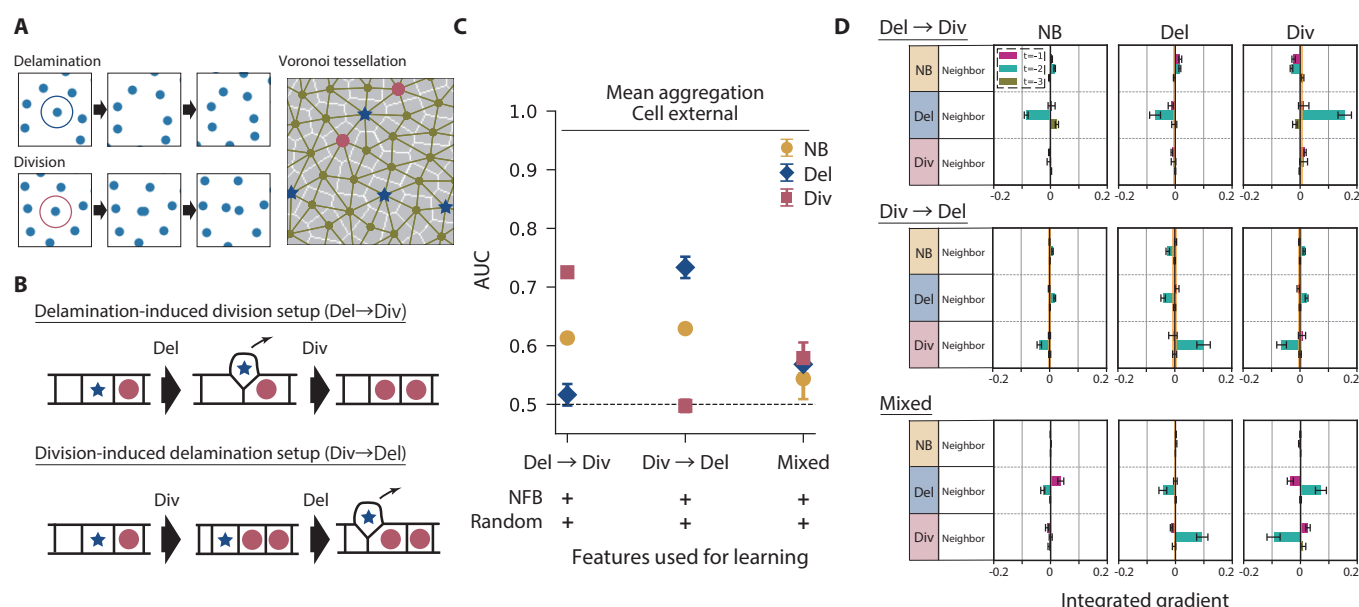


Fig. 4. Validation of the GNN approach by *in silico* models of homeostasis. (A) Snapshots of simulation from the delamination-induced division setup. Cell contacts are defined by the Voronoi tessellation. Delaminating cells and dividing cells are represented by blue stars and red circle markers, respectively. (B) Schematic of the rules applied in *in silico* models of homeostasis: delamination-induced division, division-induced delamination, and mixed. (C) The AUC of the cell external model with mean aggregation is shown for the three *in silico* setups, obtained by averaging over six trained models. Error bar: standard deviation. (D) The attribution of the cell external model with mean aggregation is shown for three *in silico* setups. The IG averaged over six trained models is shown for each pooled feature. Error bar: standard error. The upper and lower values of the IG of the random feature are shown as the orange zone.

that the AUC scores of each fate are in the expected order (Fig. 4C); in the delamination-induced division setup, the division was predictable and delamination was not, and vice versa. By testing the attribution method on these simulated data, we confirmed that the inferred mechanisms indicate the implemented rules (Fig. 4D). We noticed that the predictability of the fates become significantly low for the mixed rule setup; even in such situation, the attribution method was able to pull out the coordination rules.

Finally, we analyzed the recently obtained mouse ear skin data [37] to compare the mechanism of homeostasis across different tissues. Although the ear skin has more structure (i.e., hair follicles and other appendages), we have previously observed that the behaviors of the interfollicular epidermal basal cells in the ear skin are similar to that of the hairless hind paw [14, 34]. A noticeable difference is that the rate of divisions and delaminations were approximately two-fold slower in the ear, which is why the interval between the frames was set as 24 hours for the ear data collection. In Fig. 5, we show the AUC and the attributions for the ear data; the largest positive (negative) IG score for Div (Del) is attributed to the Del of neighboring cells from one frame earlier, indicating the existence of the same rule as the hind paw epidermis.

In addition, there are positively (negatively) large IG scores for Del (Div) attributed to the Div of neighboring cells, which were not seen in the hind paw epidermis. These scores imply that a mechanism of neighboring Div-induced Del and

Div-suppression may exist; the rule of homeostasis in the ear epidermis may be closer to the mixed setup in the simulation compared with the hind paw epidermis. This difference was confirmed by conducting a neighbor-fate imbalance analysis (fig. S8A,B). In this analysis, we focus on individual cells that either divided or delaminated and follow the subsequent behaviors of their six-nearest neighbor cells to calculate the neighbor fate imbalance [14]. The average fate imbalance indeed deviates from zero for the neighbors of dividing cells in the ear, similar to the case of the simulation with the mixed rule setup (fig. S8C-E).

DISCUSSION

Here we have used the graph-based learning framework to systematically infer how interactions correlate with cell fates in the basal layer stem cell pool of the epidermis. The GNN models were able to predict the future cell fate using the relationships of the target cell and the cells in contact in the past time frames. We identified the cell features that are attributed to the predictability of the model using IG, and clarified that neighboring cell fates have a delayed effect on the target cell's fate outcome. The GNN-based approach gives interpretable attributions in the words of the spatiotemporal relationship between the cells, which is an advantage compared with direct image-based approaches employing CNNs.

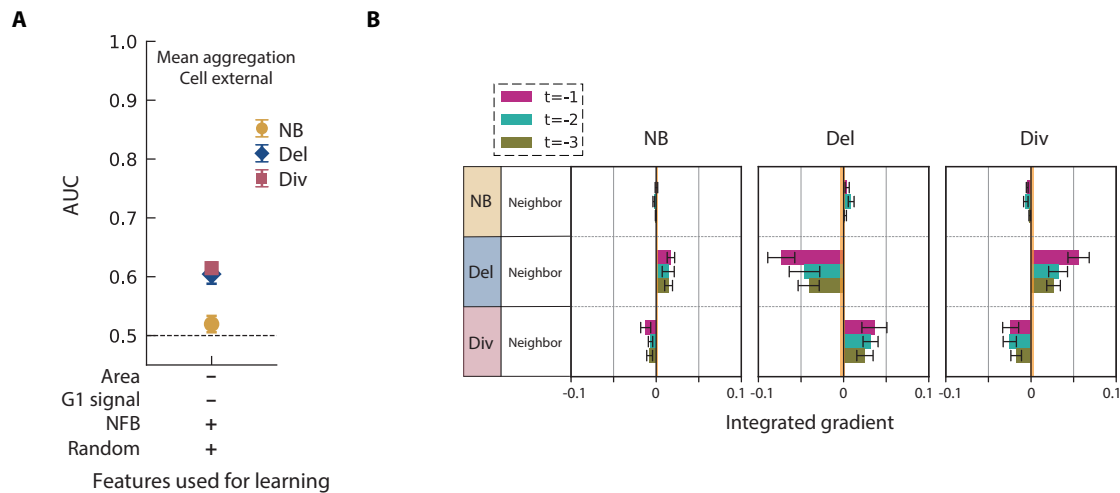


Fig. 5. Analysis of the ear epidermis data. (A) The AUC of the cell external model with mean aggregation for the ear epidermis data, obtained by averaging over six trained models. Error bar: standard deviation. (B) The attribution of the cell external model with mean aggregation is shown for ear data for the feature set: (Area, G1 signal, NFB, Random)=(−, −, +, +). The IG averaged over six trained models is shown for each pooled feature. Error bar: standard error. The upper and lower values of the IG of the random feature are shown as the orange zone.

The extracted rules from the dynamics in the epidermis included the previously reported delamination-induced neighboring division [14] as well as potentially novel interactions: suppression (induction) of cell delamination by neighboring delamination (division). Inhibition of certain behaviors by cells in contact is a common motif observed in cell biology [38–40]. Neighboring division-induced delamination may be associated with local mechanical forces [41]. Testing whether known molecular pathways responsible for lateral inhibition and mechanical signaling are playing roles in tissue homeostasis will be an interesting next step.

The predicted rules are favorable mechanisms in keeping the cell density in the homeostatic tissue, and it is intriguing that there exist differences in the rules adopted in distinct regions of the skin. A possible explanation for this is the cell density; the hind paw epidermis is crowded compared with the ear (26,000 cells/mm² versus 14,000 cells/mm², with almost the same thickness, 15 μ m, fig. S9). In high cell density regions such as the hind paw, it is reasonable that spontaneous cell division is suppressed and cells can only proliferate when there is space provided by neighboring delamination. Interestingly, in the developing skin where the cell density is much lower than the adult ear, the stratification of cells has been reported to be driven by neighboring cell division [42]. These results imply how the balance between cell differentiation and proliferation can be maintained through distinct mechanisms depending on the operating cell density regimes.

In predicting the kinetics, multiple time frames were required as inputs in our example (fig. S7A) due to the significant time delay between the cell fate events. This indicates that the proper treatment of the temporal axis was necessary, which is why GNN approaches based on static graphs [43, 44] are not directly applicable to the problem of multicellular kinetics. Furthermore, although different models have recently

been proposed to deal with evolving graphs with nodes and edges being added or removed [45–47], to the best of our knowledge, there has been no example of a graph-based approach dealing with replication of nodes (i.e., cell division). Since the incorporation of the cell tracks including cell division is fundamental in the analysis of multi-cellular systems, this study serves as an important step towards building an unbiased rule extraction framework for tissue dynamics.

Multicellular dynamics is inherently stochastic owing to the complex interaction between cells and the environment as well as the single-cell level fluctuations. Likely due to this stochasticity, the prediction scores generated by the GNN in this study were far from accurate. Nevertheless, the models were able to provide sensible predictions and cell interaction rules from the epidermis data and inferred the correct rules from the simulation data. The success of the current approach on data generated by highly-stochastic kinetics suggests that the framework is applicable not only for general multi-cellular dynamics such as developing embryos, malignant tissues, and organoids, but also for a wide range of systems where modeling by stochastic interacting agents is effective such as in disease spreading and ecology.

MATERIALS AND METHODS

Data preparation

The cell track data of the mouse hind paw epidermis basal layer was generated in our previous work [14].

For the ear epidermis data, we used the images collected in the work of [37], and conducted the semi-automated tracking procedure similar to the previous method [14]. We first performed 3D segmentation by cellpose [23] using the nucleus

channel (K14H2BmCherry) from a region size of $0.3 \text{ mm} \times 0.3 \text{ mm} \times 40 \text{ }\mu\text{m}$ and obtained the 3D masks of the cell nuclei. We then defined the height of the interface between the epidermis and the dermis based on the 3D masks of the nuclei and subtracted this height from the original 3D data to level the basal layer position. From the height-corrected 3D images, we took three consecutive z-positions containing the nucleus of all the basal layer cells and averaged the intensity over the three slices to obtain 2D images in each channel. We calculated the local maxima of the cell nuclei (K14H2BmCherry) and automatically corrected the shifts between time frames by minimizing the square distance between the nearest cell positions across the frames using affine transformation. The cells included in all ten time frames were used in the following analysis. The crop size in the 2D plane therefore varied across different areas, ranging from $139 \text{ }\mu\text{m} \times 139 \text{ }\mu\text{m}$ to $238 \text{ }\mu\text{m} \times 238 \text{ }\mu\text{m}$.

At each time frame, we segmented the cells using the marker-controlled watershed algorithm using the mem-tTomato channel and the maxima of the cell nuclei positions. We assigned each cell (i.e., segmented area) to a segmented area in the previous time frame with the largest overlap. Tracked cells were frequently lost or were associated with more than one cell in the subsequent time frame, which indicated cell delamination from the basal layer and cell division, respectively. We then manually corrected the errors in the tracking with a guide from the height-corrected 3D images using a pipeline employing napari [48], and cropped out the region close to the edge so that the remaining region only includes reliably tracked cells. The script outputted the segmented areas of the cells at each time point and their lineages, which was used to build the spatiotemporal graphs. All codes involved in generating the basal layer cell tracks in the ear epidermis were written in Python.

We used two areas of tracking data for the hind paw data, and six areas for the ear data. The total number of cell-frames in the spatiotemporal graphs was 5,996 for the hind paw data (~ 214 unique cells per area per time frame) and 12,828 for the ear data (~ 267 unique cells per area per time frame). We split the graph data into two, the training set and the test set. The number of cell-frames in the two sets was made to be comparable. The proportion of cells that experience delamination and division in each frame were both 9 % in the hind paw epidermis data and both 7 % in the ear epidermis data.

For the neighbor fate net imbalance analysis, we applied the same method as previously described [14] to the ear epidermis data as well as to the simulation data.

Graph neural network

Bidirectional spatiotemporal GNN model

The spatiotemporal graphs were created from N_t sequential time frames from the segmented time-lapse images using Deep Graph Library [49]. We first created the cell-contact graphs for each time frame. The neighboring cells α and β were connected by two directed edges pointing at each other.

We then added directional edges between the same cells in the future and the past as well as their parents and daughter cells in the sequential time frames. The cell feature vector $\mathbf{x}_{i,\alpha} \in \mathbb{R}^n$ was assigned to each cell α for the i -th time frame. We used the cell area and G1-phase reporter signal as components of the feature vector, which were obtained from the original segmented images. The features were normalized by dividing the values by the maximum values among all the training and test data. When reducing a feature, we set that particular feature to zero in all nodes.

We processed these spatiotemporal graphs with a GNN model using PyTorch. Our first GNN model is a collection of models which consists of a backward temporal edge model $\Phi^{\text{B,edge}}$, a backward temporal node model $\Phi^{\text{B,node}}$, a spatial edge model ϕ^{edge} , a spatial node model ϕ^{node} , a forward temporal edge model $\Phi^{\text{F,edge}}$, a forward temporal node model $\Phi^{\text{F,node}}$, and a decoder ψ^{dec} .

First, we propagated the information from the future to the past using $\Phi^{\text{B,edge}}$, and then updated the node feature by $\Phi^{\text{B,node}}$. By initializing the node feature $\mathbf{a}_{i,\alpha}^{(0)} = \mathbf{x}_{i,\alpha}$, the $(k+1)$ -th update ($0 \leq k \leq N_t - 2$) of the node feature $\mathbf{a}_{i,\alpha}^{(k+1)}$ of cell α in the i -th frame ($0 \leq i \leq N_t - 1$) is given by,

$$\mathbf{A}_{(i,\alpha) \leftarrow (i+1,\beta)}^{(k+1)} = \Phi^{\text{B,edge}}(\mathbf{a}_{i,\alpha}^{(k)}, \mathbf{a}_{i+1,\beta}^{(k)}), \quad (1)$$

$$\mathbf{B}_{i,\alpha}^{(k+1)} = \sum_{(i+1,\beta) \in D(i,\alpha)} \mathbf{A}_{(i,\alpha) \leftarrow (i+1,\beta)}^{(k+1)}, \quad (2)$$

$$\mathbf{a}_{i,\alpha}^{(k+1)} = \Phi^{\text{B,node}}(\mathbf{B}_{i,\alpha}^{(k+1)}, \mathbf{a}_{i,\alpha}^{(k)}). \quad (3)$$

Here, the subscript (i, α) denotes cell α in the i -th frame, and $(i, \alpha) \leftarrow (j, \beta)$ denotes the edge between cell α in the i -th frame and cell β in j -th frame. We calculated Eq. 1 for all the connected pairs of cells in the spatiotemporal graphs. $D(i, \alpha)$ is the set of the daughters of the cell α in the i -th frame if the cell α divides in the i -th frame; otherwise, $D(i, \alpha)$ is cell α itself in the $(i+1)$ -th frame. We set the features of the cells in the final frame $\mathbf{a}_{i,\alpha}^{(k+1)} = \mathbf{a}_{i,\alpha}^{(0)}$, which was not updated. Furthermore, for cell α which delaminates or exits the field of view in the i -th frame, we set $\mathbf{B}_{i,\alpha}^{(k+1)} = 0$.

Second, we calculated the edge features in each time frame using ϕ^{edge} , and then updated the cell feature vectors with the edge features using ϕ^{node} . By initializing with $\mathbf{b}_{i,\alpha}^{(0)} = \mathbf{a}_{i,\alpha}^{(N_t-1)}$, the $(l+1)$ -th update of the node feature $\mathbf{b}_{i,\alpha}^{(l+1)}$ of cell α in the i -th frame is given by,

$$\mathbf{H}_{(i,\alpha) \leftarrow (i,\beta)}^{(l+1)} = \phi^{\text{edge}}(\mathbf{b}_{i,\alpha}^{(l)}, \mathbf{b}_{i,\beta}^{(l)}), \quad (4)$$

$$\mathbf{I}_{i,\alpha}^{(l+1)} = \text{AGG}(\{\mathbf{H}_{(i,\alpha) \leftarrow (i,\beta)}^{(l+1)} : (i,\beta) \in N(i,\alpha)\}) \quad (5)$$

$$\mathbf{b}_{i,\alpha}^{(l+1)} = \phi^{\text{node}}(\mathbf{I}_{i,\alpha}^{(l+1)}, \mathbf{b}_{i,\alpha}^{(l)}). \quad (6)$$

Here, $N(i, \alpha)$ is the set of the neighbor cells of α in the i -th frame. In Eq. 5, AGG represents either the sum aggregation or the mean aggregation across the set of the neighbor cells. We repeat this process for N_s times to take into account the N_s -step neighbor interactions, in which case $0 \leq l \leq N_s - 1$.

Third, we propagated the information from the past to the future using $\Phi^{\text{F,edge}}$, and then updated the node feature by $\Phi^{\text{F,node}}$. By initializing with $\mathbf{c}_{i,\alpha}^{(0)} = \mathbf{b}_{i,\alpha}^{(N_s)}$, the $(m+1)$ -th update ($0 \leq m \leq N_t - 2$) of the node feature $\mathbf{c}_{i,\alpha}^{(m+1)}$ of cell α in

the i -th frame is given by,

$$U_{(i+1,\alpha) \in (i,\beta)}^{(m+1)} = \Phi^{\text{F,edge}}(c_{i,\alpha}^{(m)}, c_{i+1,\beta}^{(m)}), \quad (7)$$

$$V_{i,\alpha}^{(m+1)} = U_{(i+1,\alpha) \in P(i+1,\alpha)}^{(m+1)}, \quad (8)$$

$$c_{i,\alpha}^{(m+1)} = \Phi^{\text{F,node}}(V_{i,\alpha}^{(m+1)}, c_{i,\alpha}^{(m)}). \quad (9)$$

Here, $P(i+1, \alpha)$ is the parent of cell α in the $(i+1)$ -th frame if the cell α is born in the $(i+1)$ -th frame; otherwise, $P(i+1, \alpha)$ is cell α itself in the i -th frame. We set the features of the cells in the final frame $c_{i,\alpha}^{(m+1)} = c_{i,\alpha}^{(0)}$, which was not updated. Furthermore, for cell α which pops in the field of view from the outside in the i -th frame, we set $V_{i,\alpha}^{(m+1)} = 0$.

Finally, we decoded the cell feature of cell α in the final frame $((N_t - 1)\text{th frame})$ by:

$$Y_{N_t-1,\alpha} = \psi^{\text{dec}}(c_{N_t-1,\alpha}^{(N_t-1)}). \quad (10)$$

For all the functions $\Phi^{\text{B,edge}} : \mathbb{R}^{2n} \rightarrow \mathbb{R}^n$, $\Phi^{\text{B,node}} : \mathbb{R}^{2n} \rightarrow \mathbb{R}^n$, $\phi^{\text{edge}} : \mathbb{R}^{2n} \rightarrow \mathbb{R}^n$, $\phi^{\text{node}} : \mathbb{R}^{2n} \rightarrow \mathbb{R}^n$, $\Phi^{\text{F,edge}} : \mathbb{R}^{2n} \rightarrow \mathbb{R}^n$, $\Phi^{\text{F,node}} : \mathbb{R}^{2n} \rightarrow \mathbb{R}^n$, $\psi^{\text{dec}} : \mathbb{R}^n \rightarrow \mathbb{R}^3$ in our GNN model, we used the multi-layer perceptron (MLP), whose components are (1) N_{layer} hidden layers which are respectively composed of a fully-connected layer and a rectified linear unit (ReLU), and (2) an output fully-connected layer.

Unidirectional spatiotemporal GNN model

For the unidirectional spatiotemporal model, we skipped the backward temporal edge and node models in the bidirectional spatiotemporal model. Hence, we initialized $b_{i,\alpha}^{(0)} = x_{i,\alpha}$ in the spatial edge and node models. In the cell external model, we assigned a null feature vector for the target cells; in the spatial edge model, we initialized $b_{i,\alpha}^{(0)} = 0$ and $b_{i,\beta}^{(0)} = x_{i,\beta}$ for the edges from cell β to cell α in the i -th frame.

To represent the local lineage branch structure, we introduced the next frame behavior (NFB) as a new feature, which encodes the behavior of the next frame of that cell by a one-hot vector, NB $([1, 0, 0])$, Del $([0, 1, 0])$, or Div $([0, 0, 1])$. Since the NFB in the final frame of each network is what we aim to predict, we set the NFB in the final frame to null vector $([0, 0, 0])$.

Training

Since the proportion of the three cell fates is imbalanced, we used the weighted softmax-cross-entropy loss where the weight of each label was set to the inverse of the proportion of the cell fate in the training data. To minimize the loss, we used Adam optimizer with the learning rate $lr = 0.0001$. In the training, we input a spatiotemporal graph of N_t sequential time frames obtained from an imaging area to update the parameters of the GNN model. We repeated the update for all the spatiotemporal graphs in a single epoch of the training.

To optimize the number of layers N_{layer} , the number of nodes N_{node} of a hidden layer of the MLPs, and the dropout

rate p , we tested the performance of the GNN model by changing these parameters. For this test, we used the simulation data of delamination-induced division setup with NFB and random features. First, we investigated the effect of N_{layer} by setting $N_{\text{node}} = 50$ and $p = 0$. To quantify the performance of the GNN model, we calculated the maximum value of macro-F1 score of the test data during learning for 2000 epochs. We ran 5000 epochs for $N_{\text{layer}} = 2$, since the learning was exceptionally slow. As shown in fig. S10A, $N_{\text{layer}} = 1$ and $N_{\text{layer}} = 2$ give comparable performance. Second, we investigated the effect of N_{node} by setting $N_{\text{layer}} = 1$ and $p = 0$. We found that N_{node} does not significantly affect the performance (fig. S10B). Hence, we chose $N_{\text{layer}} = 1$ and $N_{\text{node}} = 50$ to minimize the size of the GNN model. Finally, we changed p as shown in fig. S10C, and found that p also does not significantly affect the performance. We here chose $p = 0.1$ since the average performance was slightly better and the standard deviation is smaller compared with the other conditions. In all the training, we set $N_{\text{layer}} = 1$, $N_{\text{node}} = 50$ and $p = 0.1$, and ran 2000 epochs. We exceptionally ran 5000 epochs for the bidirectional spatiotemporal model with mean aggregation without any feature.

We also tested the effect of data size on the performance (fig. S11). We found that the prediction score and attribution score are sufficiently high when the number of cell-frames used in the training data is above 2000.

The AUC and attribution of each condition were calculated for the GNN model obtained at the epoch at which the model achieves the maximum macro-F1 score [50]. In fig. S2, we show an example of the learning curves of the cell external model for the hind paw data with NFB and the random feature. The weighted softmax-cross-entropy loss, macro-F1 score, recall, and precision curves are shown respectively in fig. S2. The vertical lines in figs. S2B-D indicate the epoch at which the model achieves the best macro-F1 score.

Attribution method

We used the integrated gradients (IG, [36]) for the attribution. The IG $I_k(g)$ of the k -th feature for an input subgraph g of a target cell is given by:

$$I_k^f(g) = (X_k - X'_k) \cdot \int_0^1 \frac{\partial F^f(\mathbf{X}' + \alpha \cdot (\mathbf{X} - \mathbf{X}'))}{\partial X_k} d\alpha. \quad (11)$$

The function F^f is the softmax score of fate $f \in \{\text{NB}, \text{Del}, \text{Div}\}$ as a function of the input features calculated by the trained network. \mathbf{X} represents the concatenated features of all the cells of g , and X_k is the value of the k th feature. \mathbf{X}' is the baseline, which is the null vector of the same size as \mathbf{X} . We calculated the IG for three cell fate labels of all the input graphs.

Since the baseline should be neutral for calculating the attribution, the GNN model must be trained for the null graphs to provide equal soft-max scores for three cell fate labels. To this end, we minimized the mean-squared-error (MSE) loss

defined as,

$$L_{\text{MSE}} = \frac{1}{N_{\text{sub}}} \sum_f \sum_{n=1}^{N_{\text{sub}}} \left(p_n^{f, \text{null}} - \frac{1}{3} \right)^2 \quad (12)$$

together with the weighted softmax-cross-entropy loss. Here, $p_i^{f, \text{null}}$ is the softmax score of fate $f \in \{\text{NB}, \text{Del}, \text{Div}\}$ of the null graph of the n -th target cell's subgraph, and N_{sub} is the number of the subgraphs. Within a single epoch, we first input a spatiotemporal graph to update the parameters using the weighted softmax-cross-entropy loss, and then input the corresponding null graph to update the parameters using the MSE loss. By this learning method, we approximately obtained the neutral baseline softmax scores for null graphs (fig. S12).

In the analysis of the attribution, we pooled each feature of each target cell's subgraph into the relative spatiotemporal position of the feature with respect to the target cell. We calculated the average IG of each category for each cell fate label. With respect to the average IG of NFB, we averaged the IG of each category of NFB only among the cells in the category. Finally, we calculated the mean of the average IG over all the trained GNN models. We also defined the baseline for non-zero signals shown as the orange shades in the attribution plots. The minimum (maximum) value of the range of the baseline is defined as the minimum (maximum) of the model-average of IG subtracted (added) by the standard error among all the pooled random features.

Since $\sum_f F^f(\mathbf{X}) = 1$ for any \mathbf{X} , we have

$$\sum_f I_k^f(\mathbf{g}) = 0, \quad (13)$$

To make the plots in Figs. 2-5, we pooled the features according to the spatiotemporal position. The pooled features, which we denote as a , contain multiple features k in the original calculation. The attribution score for each pooled feature is evaluated by

$$I_a^f(\mathbf{g}) = \frac{\sum_{k \in a} I_k^f(\mathbf{g})}{\sum_{k \in a} 1}. \quad (14)$$

Note that the normalization still holds:

$$\sum_f I_a^f(\mathbf{g}) = \frac{\sum_{k \in a} \sum_f I_k^f(\mathbf{g})}{\sum_{k \in a} 1} = 0. \quad (15)$$

Numerical simulations

For the numerical data that mimics the dynamics of basal layer dynamics, we took a simplified model of interacting particles that exclude each other through mechanical interactions and undergo stochastic division and elimination events. We placed N_0 cells labeled by α in a two-dimensional plane with size $L \times L$ and with periodic boundary conditions and let the cells interact with each other through an interacting potential. The equation of motion reads

$$\dot{\mathbf{r}}_\alpha(t) = - \sum_{\beta \neq \alpha} \frac{\partial}{\partial \mathbf{r}_\alpha} u(\mathbf{r}_\alpha, \mathbf{r}_\beta), \quad (16)$$

which is an overdamped kinetics without noise. \mathbf{r}_α is the position vector of cell α . The repulsive interacting potential has the typical length scale l :

$$u(\mathbf{r}_\alpha, \mathbf{r}_\beta) = \begin{cases} \frac{1}{2} K (|\mathbf{r}_\alpha - \mathbf{r}_\beta| - l)^2 & |\mathbf{r}_\alpha - \mathbf{r}_\beta| < l \\ 0 & |\mathbf{r}_\alpha - \mathbf{r}_\beta| \geq l \end{cases}. \quad (17)$$

The rules of cell division and delamination were implemented by the Monte Carlo method. At the time frame of cell division of cell α , a newly born cell α' is generated at a random position within a small distance $d = 0.001 \times L$. Cell delamination is conducted by eliminating a particle instantaneously. In both cases, the position of the cells quickly relaxes to a dispersed state due to the repulsive force between the cells, Eq. 16.

The stochastic rules of fate coordination that we tested are:

- **Delamination-induced division:** delaminating cells are chosen randomly with rate λ from the pool of cells that have not yet committed to delaminate or to divide. The chosen cells are committed to delaminate and are assigned a remaining lifetime chosen from a uniform distribution between 32.4 hours and 39.6 hours. At the time point of delamination, one of the six-nearest neighbors of the delaminating cell is randomly chosen, again excluding the cells that have already committed to division or delamination, and is assigned to divide after a randomly chosen remaining lifetime drawn from a uniform distribution between 44.4 hours and 51.6 hours. $N_0 = 612$.
- **Division-induced delamination:** dividing cells are chosen randomly with the rate λ from the pool of cells that have not yet committed to delaminate or to divide. The chosen cells are committed to dividing and are assigned a remaining lifetime randomly chosen from a uniform distribution between 32.4 hours and 39.6 hours. At the time point of division, one of the six-nearest neighbors of the dividing cell is randomly chosen, again excluding the cells that have already committed to division or delamination, and is assigned to delaminate after a randomly chosen remaining lifetime drawn from a uniform distribution between 44.4 hours and 51.6 hours. $N_0 = 412$.
- **Mixed:** the two schemes explained above were mixed, with the rate of randomly assigning the delaminating and dividing cells being almost halved so that the overall event rate does not change. $N_0 = 512$.

We used $K = 9 \text{ hours}^{-1}$, $l = 0.125$, and $L = 1$ in all the simulations. For the time steps, we took $\Delta t = 1.2$ hours.

In generating the data for the graph construction, we first prepared N_0 points randomly placed inside the box (size $L \times L$) and simulated the time evolution according to Eq. 16 by the Euler method for 100 steps to obtain a dispersed cell configuration. Next, we ran the simulation up to 300 steps (360 hours) with both the equation of motion Eq. 16 and the stochastic divisions and delaminations, to make sure that the system has

reached a steady-state (fig. S13). Finally, we ran the simulation for another 300 steps and sub-sampled the time points every 20 steps (24 hours) from this final time series to generate data resembling the ear epidermis. The initial number of cells N_0 was changed in the three setups to ensure that the number of cells in the frames at steady-state are roughly the same (around 500, fig. S13). The rates of fates λ were also tuned for each setup so that the number of events that take place per cell per time frame is comparable with the experiment.

We cropped out the edges and used the data from the points in the center region $0.65 L \times 0.65 L$ so that the number of points per frame is roughly the same as the number of cells per frame in the hind paw and ear data (around 210). The neighboring cell network was generated by the two-dimensional Voronoi tessellation, and by whole spatiotemporal network was fed to the GNN learning process in the same way as the experimental data.

- [1] A. Regev, *et al.*, The Human Cell Atlas. *eLife* **6**, e27041 (2017).
- [2] K. L. Frieda, J. M. Linton, S. Hormoz, J. Choi, K.-H. K. Chow, Z. S. Singer, M. W. Budde, M. B. Elowitz, L. Cai, Synthetic recording and in situ readout of lineage information in single cells. *Nature* **541**, 107–111 (2017).
- [3] D. E. Wagner, A. M. Klein, Lineage tracing meets single-cell omics: opportunities and challenges. *Nat. Rev. Genet.* **21**, 410–427 (2020).
- [4] Y. Zhang, F. Liu, Multidimensional Single-Cell Analyses in Organ Development and Maintenance. *Trends Cell Biol.* **29**, 477–486 (2019).
- [5] G. Pasqual, A. Chudnovskiy, J. M. J. Tas, M. Agudelo, L. D. Schweitzer, A. Cui, N. Hacohen, G. D. Vitoria, Monitoring T cell–dendritic cell interactions in vivo by intercellular enzymatic labelling. *Nature* **553**, 496–500 (2018).
- [6] J.-C. Boisset, J. Vivié, D. Grün, M. J. Muraro, A. Lyubimova, A. van Oudenaarden, Mapping the physical network of cellular interactions. *Nat. Methods* **15**, 547–553 (2018).
- [7] K. B. Halpern, R. Shenhav, H. Massalha, B. Toth, A. Egozi, E. E. Massasa, C. Medgalia, E. David, A. Giladi, A. E. Moor, Z. Porat, I. Amit, S. Itzkovitz, Paired-cell sequencing enables spatial gene expression mapping of liver endothelial cells. *Nat. Biotechnol.* **36**, 962–970 (2018).
- [8] R. Tang, C. W. Murray, I. L. Linde, N. J. Kramer, Z. Lyu, M. K. Tsai, L. C. Chen, H. Cai, A. D. Gitler, E. Engleman, W. Lee, M. M. Winslow, A versatile system to record cell-cell interactions. *eLife* **9**, e61080 (2020).
- [9] A. Giladi, M. Cohen, C. Medaglia, Y. Baran, B. Li, M. Zada, P. Bost, R. Blecher-Gonen, T.-M. Salame, J. U. Mayer, E. David, F. Ronchese, A. Tanay, I. Amit, Dissecting cellular crosstalk by sequencing physically interacting cells. *Nat. Biotechnol.* **38**, 629–637 (2020).
- [10] C. Baccin, J. Al-Sabah, L. Velten, P. M. Helbling, F. Grünschlager, P. Hernández-Malmierca, C. Nombela-Arrieta, L. M. Steinmetz, A. Trumpp, S. Haas, Combined single-cell and spatial transcriptomics reveal the molecular, cellular and spatial bone marrow niche organization. *Nat. Cell Biol.* **22**, 38–48 (2020).
- [11] I. C. Clark, C. Gutiérrez-Vázquez, M. A. Wheeler, Z. Li, V. Rothhammer, M. Linnerbauer, L. M. Sanmarco, L. Guo, M. Blain, S. E. Zandee, *et al.*, Barcoded viral tracing of single-cell interactions in central nervous system inflammation. *Science* **372**, eabf1230 (2021).
- [12] E. Clayton, D. P. Doupe, A. M. Klein, D. J. Winton, B. D. Simons, P. H. Jones, A single type of progenitor cell maintains normal epidermis. *Nature* **446**, 185–189 (2007).
- [13] A. M. Klein, B. D. Simons, Universal patterns of stem cell fate in cycling adult tissues. *Development* **138**, 3103–3111 (2011).
- [14] K. R. Mesa, K. Kawaguchi, K. Cockburn, D. Gonzalez, J. Boucher, T. Xin, A. M. Klein, V. Greco, Homeostatic Epidermal Stem Cell Self-Renewal Is Driven by Local Differentiation. *Cell Stem Cell* **23**, 677–686.e4 (2018).
- [15] J. Liang, S. Balachandra, S. Ngo, L. E. O’Brien, Feedback regulation of steady-state epithelial turnover and organ size. *Nature* **548**, 588–591 (2017).
- [16] L. Ritsma, S. I. Ellenbroek, A. Zomer, H. J. Snippert, F. J. de Sauvage, B. D. Simons, H. Clevers, J. van Rheenen, Intestinal crypt homeostasis revealed at single-stem-cell level by in vivo live imaging. *Nature* **507**, 362–365 (2014).
- [17] S. Han, J. Fink, D. J. Jörg, E. Lee, M. K. Yum, L. Chatzeli, S. R. Merker, M. Jossierand, T. Trendafilova, A. Andersson-Rolf, C. Dabrowska, H. Kim, R. Naumann, J.-H. Lee, N. Sasaki, R. L. Mort, O. Basak, H. Clevers, D. E. Stange, A. Philpott, J. K. Kim, B. D. Simons, B.-K. Koo, Defining the Identity and Dynamics of Adult Gastric Isthmus Stem Cells. *Cell Stem Cell* **25**, 342–356.e7 (2019).
- [18] H. Yamaguchi, K. Kawaguchi, T. Sagawa, Dynamical crossover in a stochastic model of cell fate decision. *Phys. Rev. E* **96** (2017).
- [19] R. R. Stine, E. L. Matunis, Stem cell competition: finding balance in the niche. *Trends Cell Biol.* **23**, 357–364 (2013).
- [20] Y. Kitadate, D. J. Jörg, M. Tokue, A. Maruyama, R. Ichikawa, S. Tsuchiya, E. Segi-Nishida, T. Nakagawa, A. Uchida, C. Kimura-Yoshida, S. Mizuno, F. Sugiyama, T. Azami, M. Ema, C. Noda, S. Kobayashi, I. Matsuo, Y. Kanai, T. Nagasawa, Y. Sugimoto, S. Takahashi, B. D. Simons, S. Yoshida, Competition for Mitogens Regulates Spermatogenic Stem Cell Homeostasis in an Open Niche. *Cell Stem Cell* **24**, 79–92.e6 (2019).
- [21] O. Ronneberger, P. Fischer, T. Brox, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Lecture Notes in Computer Science (Springer International Publishing, Cham, 2015), pp. 234–241.
- [22] T. Falk, D. Mai, R. Bensch, O. Çiçek, A. Abdulkadir, Y. Marrakchi, A. Böhm, J. Deubner, Z. Jäckel, K. Seiwald, A. Dovzhenko, O. Tietz, C. D. Bosco, S. Walsh, D. Saltukoglu, T. L. Tay, M. Prinz, K. Palme, M. Simons, I. Diester, T. Brox, O. Ronneberger, U-Net: deep learning for cell counting, detection, and morphometry. *Nat. Methods* **16**, 67 (2019).
- [23] C. Stringer, T. Wang, M. Michaelos, M. Pachitariu, Cellpose: a generalist algorithm for cellular segmentation. *Nat. Methods* **18**, 100–106 (2021).
- [24] E. M. Christiansen, S. J. Yang, D. M. Ando, A. Javaherian, G. Skibinski, S. Lipnick, E. Mount, A. O’Neil, K. Shah, A. K. Lee, *et al.*, In silico labeling: predicting fluorescent labels in unlabeled images. *Cell* **173**, 792–803 (2018).
- [25] F. Buggenthin, F. Buettner, P. S. Hoppe, M. Ende, M. Kroiss, M. Strasser, M. Schwarzfischer, D. Loeffler, K. D. Kokkalis,

- O. Hilsenbeck, T. Schroeder, F. J. Theis, C. Marr, Prospective identification of hematopoietic lineage choice by deep learning. *Nat. Methods* **14**, 403–406 (2017).
- [26] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, G. Monfardini, The graph neural network model. *IEEE transactions on neural networks* **20**, 61–80 (2008).
- [27] P. W. Battaglia, R. Pascanu, M. Lai, D. Rezende, K. Kavukcuoglu, Interaction networks for learning about objects, relations and physics. arXiv:1612.00222 (2016).
- [28] M. Cranmer, A. Sanchez-Gonzalez, P. Battaglia, R. Xu, K. Cranmer, D. Spergel, S. Ho, Discovering Symbolic Models from Deep Learning with Inductive Biases. arXiv:2006.11287 (2020).
- [29] V. Bapst, T. Keck, A. Grabska-Barwińska, C. Donner, E. D. Cubuk, S. S. Schoenholz, A. Obika, A. W. R. Nelson, T. Back, D. Hassabis, P. Kohli, Unveiling the predictive power of static structure in glassy systems. *Nat. Phys.* **16**, 448–454 (2020).
- [30] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, U. Alon, Network motifs: simple building blocks of complex networks. *Science* **298**, 824–827 (2002).
- [31] A. Sakaue-Sawano, H. Kurokawa, T. Morimura, A. Hanyu, H. Hama, H. Osawa, S. Kashiwagi, K. Fukami, T. Miyata, H. Miyoshi, T. Imamura, M. Ogawa, H. Masai, A. Miyawaki, Visualizing Spatiotemporal Dynamics of Multicellular Cell-Cycle Progression. *Cell* **132**, 487–498 (2008).
- [32] K. Xu, W. Hu, J. Leskovec, S. Jegelka, How powerful are graph neural networks? arXiv:1810.00826 (2018).
- [33] S. Xie, J. M. Skotheim, A G1 sizer coordinates growth and division in the mouse epidermis. *Curr. Biol.* **30**, 916–924 (2020).
- [34] P. Rempoulas, K. R. Mesa, K. Kawaguchi, S. Park, D. Gonzalez, S. Brown, J. Boucher, A. M. Klein, V. Greco, Spatiotemporal coordination of stem cell commitment during epidermal homeostasis. *Science* **352**, 1471–1474 (2016).
- [35] J. A. Knoblich, Mechanisms of asymmetric stem cell division. *Cell* **132**, 583–597 (2008).
- [36] M. Sundararajan, A. Taly, Q. Yan, Axiomatic attribution for deep networks. *International Conference on Machine Learning* **70**, 3319–3328 (2017).
- [37] K. Cockburn, K. Annusver, S. Ganesan, K. R. Mesa, K. Kawaguchi, M. Kasper, V. Greco, Gradual differentiation uncoupled from cell cycle exit generates heterogeneity in the epidermal stem cell layer. *bioRxiv* 2021.01.07.425777 (2021).
- [38] R. Sancho, C. A. Cremona, A. Behrens, Stem cell and progenitor fate in the mammalian intestine: Notch and lateral inhibition in homeostasis and disease. *EMBO reports* **16**, 571–581 (2015).
- [39] H. Meinhardt, A. Gierer, Pattern formation by local self-activation and lateral inhibition. *BioEssays* **22**, 753–760 (2000).
- [40] P. A. Gagliardi, M. Dobrzyński, M.-A. Jacques, C. Dessauges, P. Ender, Y. Blum, R. M. Hughes, A. R. Cohen, O. Pertz, Collective ERK/Akt activity waves orchestrate epithelial homeostasis by driving apoptosis-induced survival. *Dev. Cell* **56**, 1–15 (2021).
- [41] G. T. Eisenhoffer, P. D. Loftus, M. Yoshigi, H. Otsuna, C.-B. Chien, P. A. Morcos, J. Rosenblatt, Crowding induces live cell extrusion to maintain homeostatic cell numbers in epithelia. *Nature* **484**, 546–549 (2012).
- [42] Y. A. Miroshnikova, H. Q. Le, D. Schneider, T. Thalheim, M. Rübsam, N. Bremicker, J. Polleux, N. Kamprad, M. Taranola, I. Wang, M. Balland, C. M. Niessen, J. Galle, S. A. Wickström, Adhesion forces and cortical tension couple cell proliferation and differentiation to drive epidermal stratification. *Nat. Cell Biol.* **20**, 69–80 (2018).
- [43] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, P. S. Yu, A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems* **32**, 4–24 (2021).
- [44] S. Georgousis, M. P. Kenning, X. Xie, Graph Deep Learning: State of the Art and Challenges. *IEEE Access* **9**, 22106–22140 (2021).
- [45] P. Goyal, N. Kamra, X. He, Y. Liu, DynGEM: Deep Embedding Method for Dynamic Graphs. arXiv:1805.11273 (2018).
- [46] A. Pareja, G. Domeniconi, J. Chen, T. Ma, T. Suzumura, H. Kanezashi, T. Kaler, T. B. Schardl, C. E. Leiserson, EvolveGCN: Evolving Graph Convolutional Networks for Dynamic Graphs. arXiv:1902.10191 (2019).
- [47] C. Wu, G. Nikolentzos, M. Vazirgiannis, EvoNet: A Neural Network for Predicting the Evolution of Dynamic Graphs. arXiv:2003.00842 (2020).
- [48] napari contributors, napari: a multi-dimensional image viewer for python. doi:10.5281/zenodo.3555620 (2019).
- [49] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao, T. He, G. Karypis, J. Li, Z. Zhang, Deep graph library: A graph-centric, highly-performant package for graph neural networks. arXiv:1909.01315 (2019).
- [50] M. Grandini, E. Bagli, G. Visani, Metrics for Multi-Class Classification: An Overview. arXiv:2008.05756 (2020).
- Acknowledgements:** We thank A. Tanaka and H. Tanaka for fruitful discussions. We also thank K. Adachi and Y. Fukai for the critical reading of the manuscript.
- Funding:** This work is supported by the funding from JSPS KAKENHI grants number 19K16096 (T.Y.), JP18H04760, JP18K13515, JP19H05275, JP19H05795, and Research Grant from Human Frontier Science Program (Ref. Grant No. RGY0081/2019) (K.K.), the New York Stem Cell Foundation, an HHMI Scholar award and NIH grants number 1R01AR063663-01, 1R01AR067755-01A1 and 1DP1AG066590-01 (V.G.), the Canadian Institutes of Health Research, and as a New York Stem Cell Foundation Druckenmiller Fellow (K.C.).
- Author Contributions:** T.Y. and K.K. conceived the project, analyzed the data, and conducted the numerical simulations. K.C. and V.G. performed experiments, collected and analyzed the data. T.Y. and K.K. wrote the manuscript with input throughout from K.C. and V.C.
- Competing Interests:** The authors declare that they have no competing interests.
- Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Additional data related to this paper may be requested from the authors. The python codes for the training of the GNN models will be available at <https://github.com/NoneqPhysLivingMatterLab>.

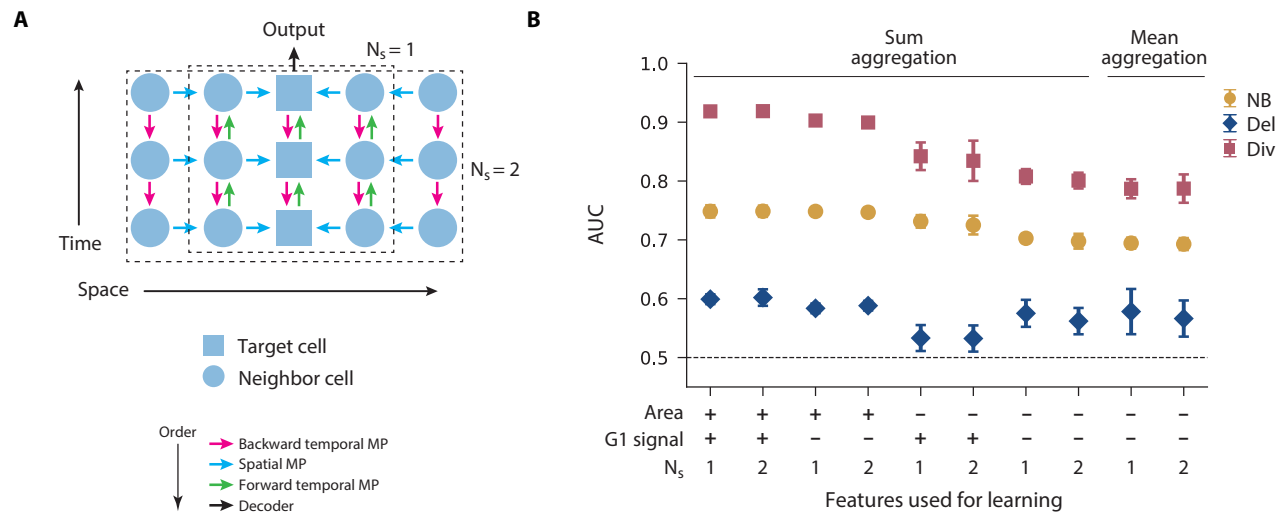


Fig. S1. Dependence of the number of iteration of the spatial MP on the performance. We apply the bidirectional spatiotemporal GNN model with sum or mean aggregation for the hind paw data under various feature conditions. **(A)** Schematic of the interaction range achieved by the different numbers of iterations of the spatial MP. **(B)** The AUC for each cell fate label obtained by averaging over six trained models. Error bar: standard deviation.

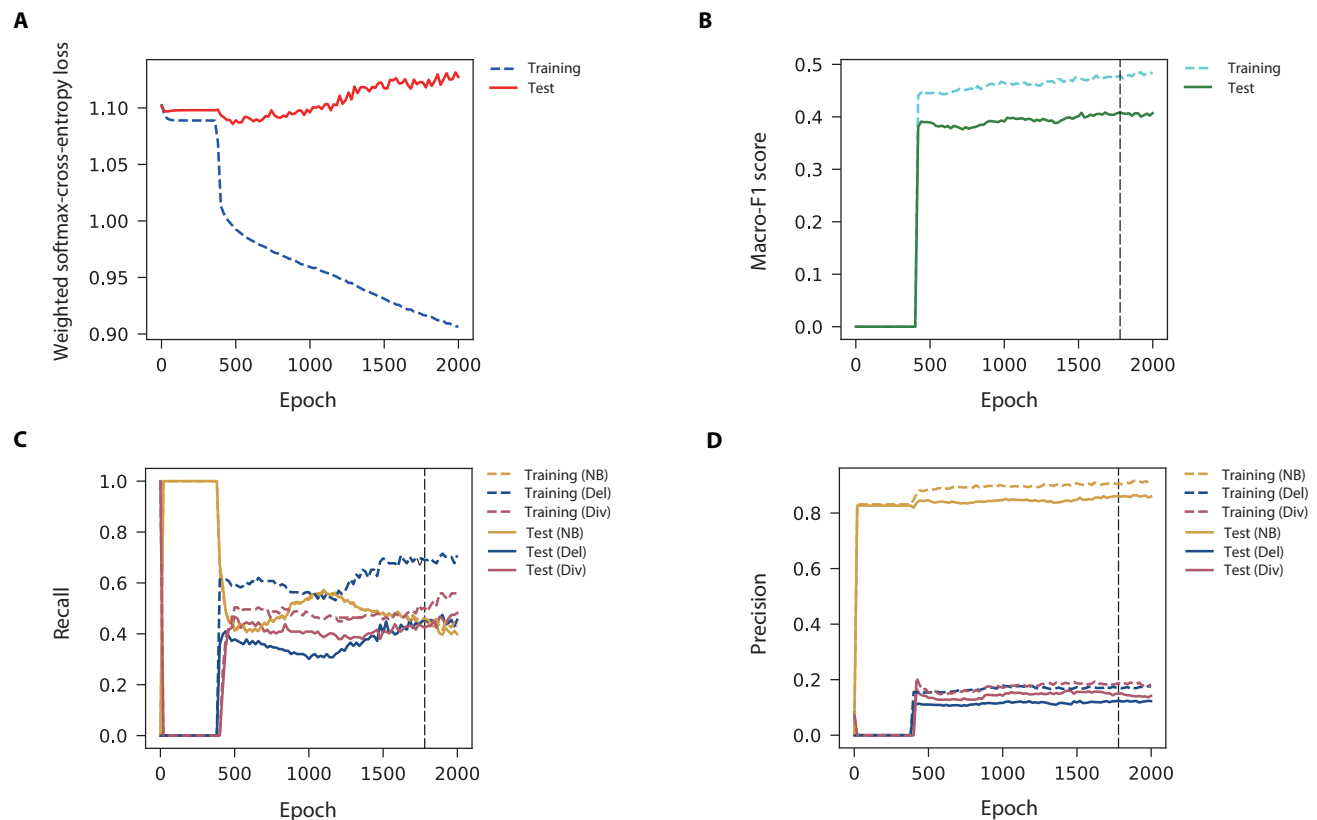


Fig. S2. Learning curves. The cell external model with mean aggregation was applied for the hind paw data with the feature condition: (Area, G1 signal, NFB, Random)=(-, -, +, +). The curves of **(A)** the weighted softmax-cross-entropy loss, **(B)** macro-F1 score, **(C)** recall and **(D)** precision for training and test data, respectively. The vertical dashed lines indicate the epoch at which the model achieves the best macro-F1 score.

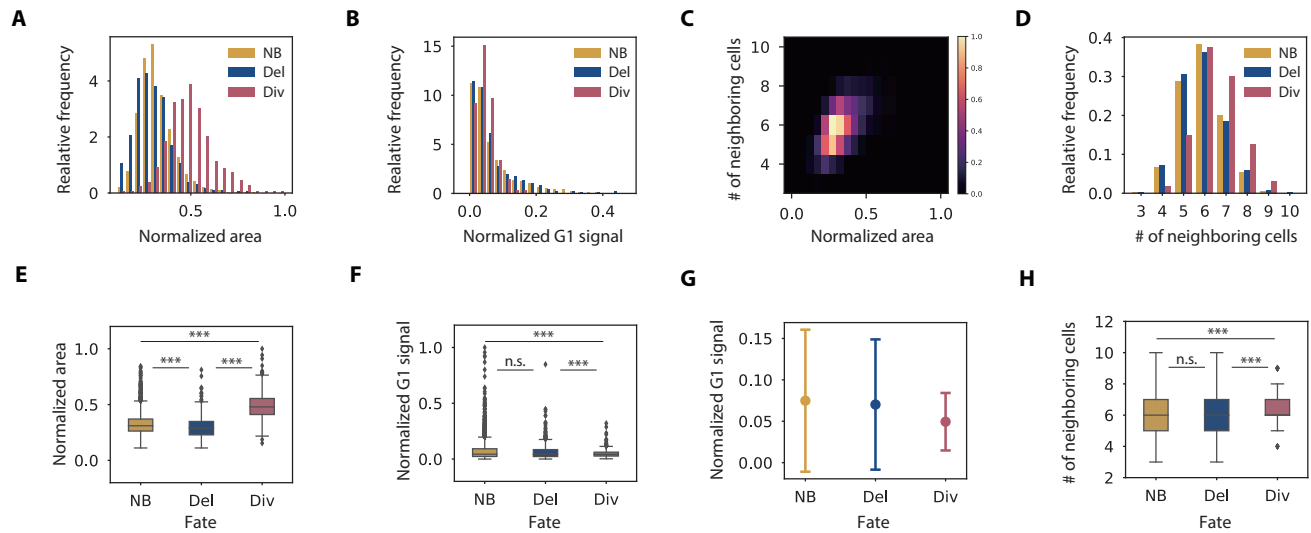


Fig. S3. **Statistical properties of the hind paw data.** (A,E) Histogram of relative frequency and the boxplot of the normalized area. (B,F,G) Histogram of relative frequency, the boxplot, and the plot of the average of the normalized G1 signal. The error bar in (G) is the standard deviation. (C) The correlation matrix between the number of neighboring cells and normalized area. (D,H) Histogram of relative frequency, the boxplot of the number of neighboring cells. In (E,F,H), the significance obtained by the Kolmogorov–Smirnov test is shown (* : $p < 0.05$, ** : $p < 0.01$, *** : $p < 0.001$).




Fig. S4. **AUC of the bidirectional spatiotemporal GNN model and the unidirectional spatiotemporal GNN model with sum aggregation.** The AUC of the bidirectional spatiotemporal GNN model (Bi.) and the unidirectional spatiotemporal GNN (Uni.) with sum aggregation are shown for various feature conditions obtained by averaging the AUC over six trained models. Error bar: standard deviation.




Fig. S5. **Effect of NFB on the performance of the unidirectional spatiotemporal GNN model with sum aggregation.** The AUC of the unidirectional spatiotemporal GNN for models under various feature sets obtained by averaging over six trained models. Error bar: standard deviation.

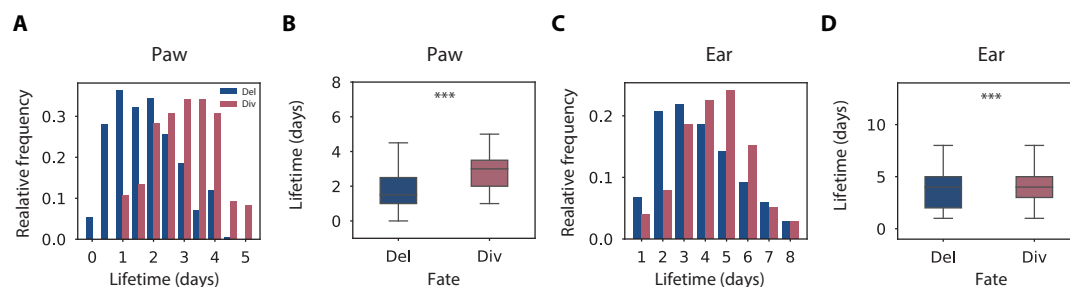


Fig. S6. **Lifetime distribution.** The lifetime distribution of delaminating and dividing cells as well as the box plots are presented for the (A,B) hind paw and (C,D) ear. In (B) and (D), the significance of Kolmogorov–Smirnov test is shown (* : $p < 0.05$, ** : $p < 0.01$, *** : $p < 0.001$).

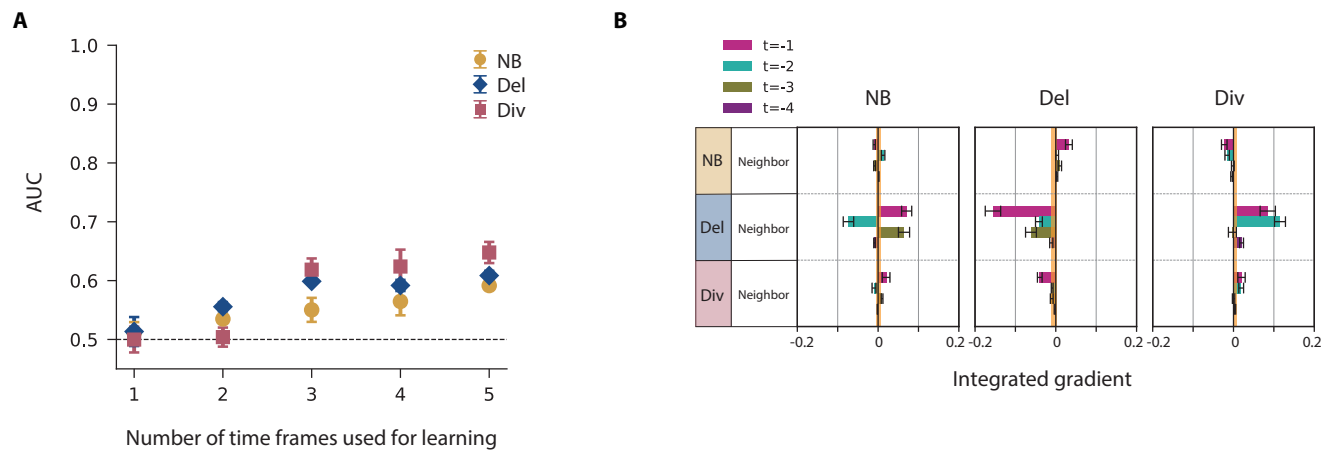


Fig. S7. **Dependence of the number of time frames on the performance.** The cell external model with mean aggregation was applied for the hind paw data with the feature condition: (NFB, Random)=(+, +). **(A)** The AUC for the different number of time frames. The AUC for each cell fate label obtained by averaging the AUC over six trained models. Error bar: standard deviation. **(B)** The attribution of the five-time model. The IG averaged over six trained models is shown for each pooled feature. Error bar: standard error.

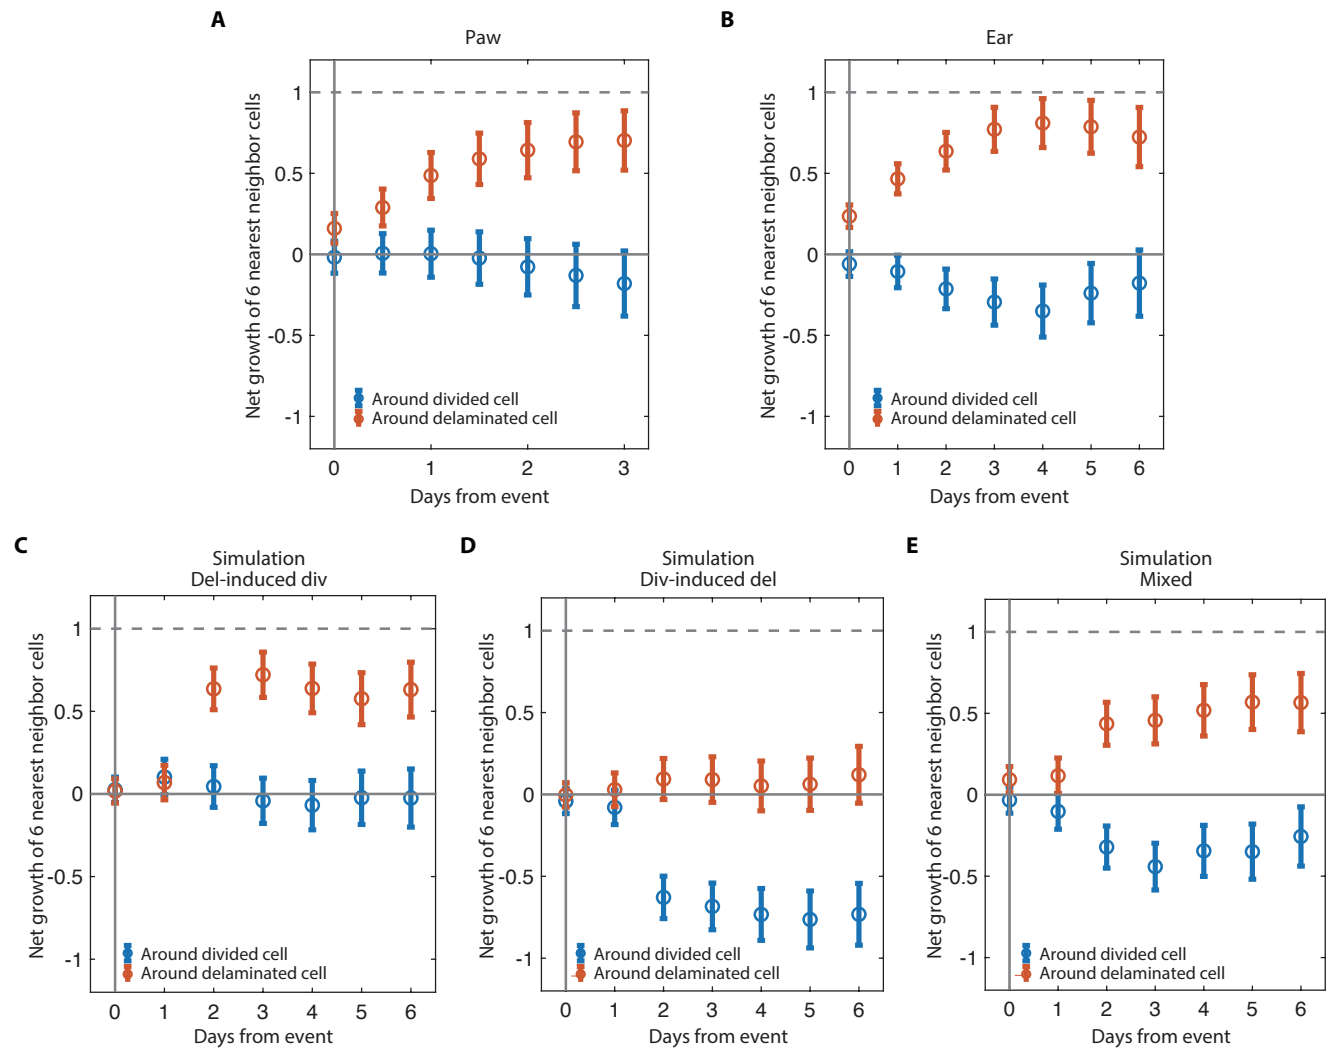


Fig. S8. Neighbor fate net imbalance analysis for the epidermis data and simulations. Net imbalance of six-nearest neighbor cells around a divided cell and a delaminated cell is plotted for the (A) hind paw and (B) ear data, as well as for data generated by simulations in the (C) delamination-induced division setup, (D) division-induced delamination setup and (E) mixed setup. Error bar: standard error.

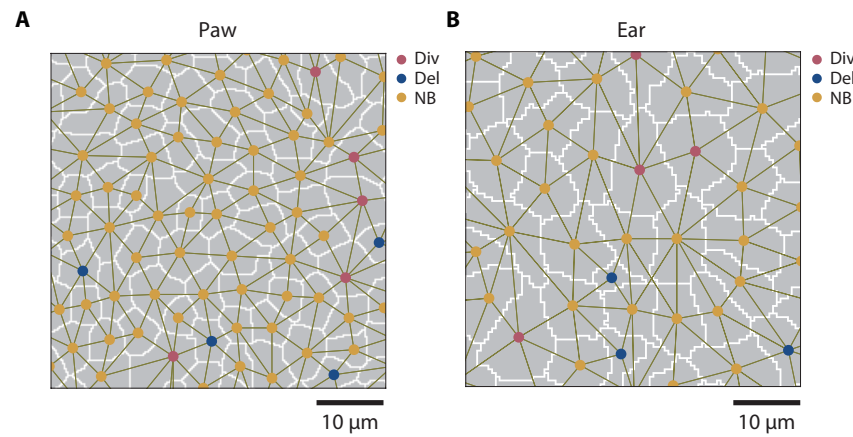


Fig. S9. **Segmented images of the hind paw and ear epidermis.** Segmented images are shown for the (A) hind paw and (B) ear epidermis. Cells undergoing Div, Del, and NB are indicated by red, blue, and yellow circle markers. The cell-contact graph is also shown.




Fig. S10. **Dependence of the performance of the GNN models on the hyperparameters.** The cell external model of the unidirectional GNN with mean aggregation is applied to the simulation data from delamination-induced division setup with the feature condition: (NFB, Random)=(+, +). The macro-F1 score averaged over six trained models is plotted against the hyperparameters. Error bar: standard deviation. (A) The effect of N_{layer} is tested with $N_{\text{node}} = 50$ and $p = 0$. (B) The effect of N_{node} is tested with $N_{\text{layer}} = 1$ and $p = 0$. (C) The effect of p is tested with $N_{\text{node}} = 50$ and $N_{\text{layer}} = 1$.

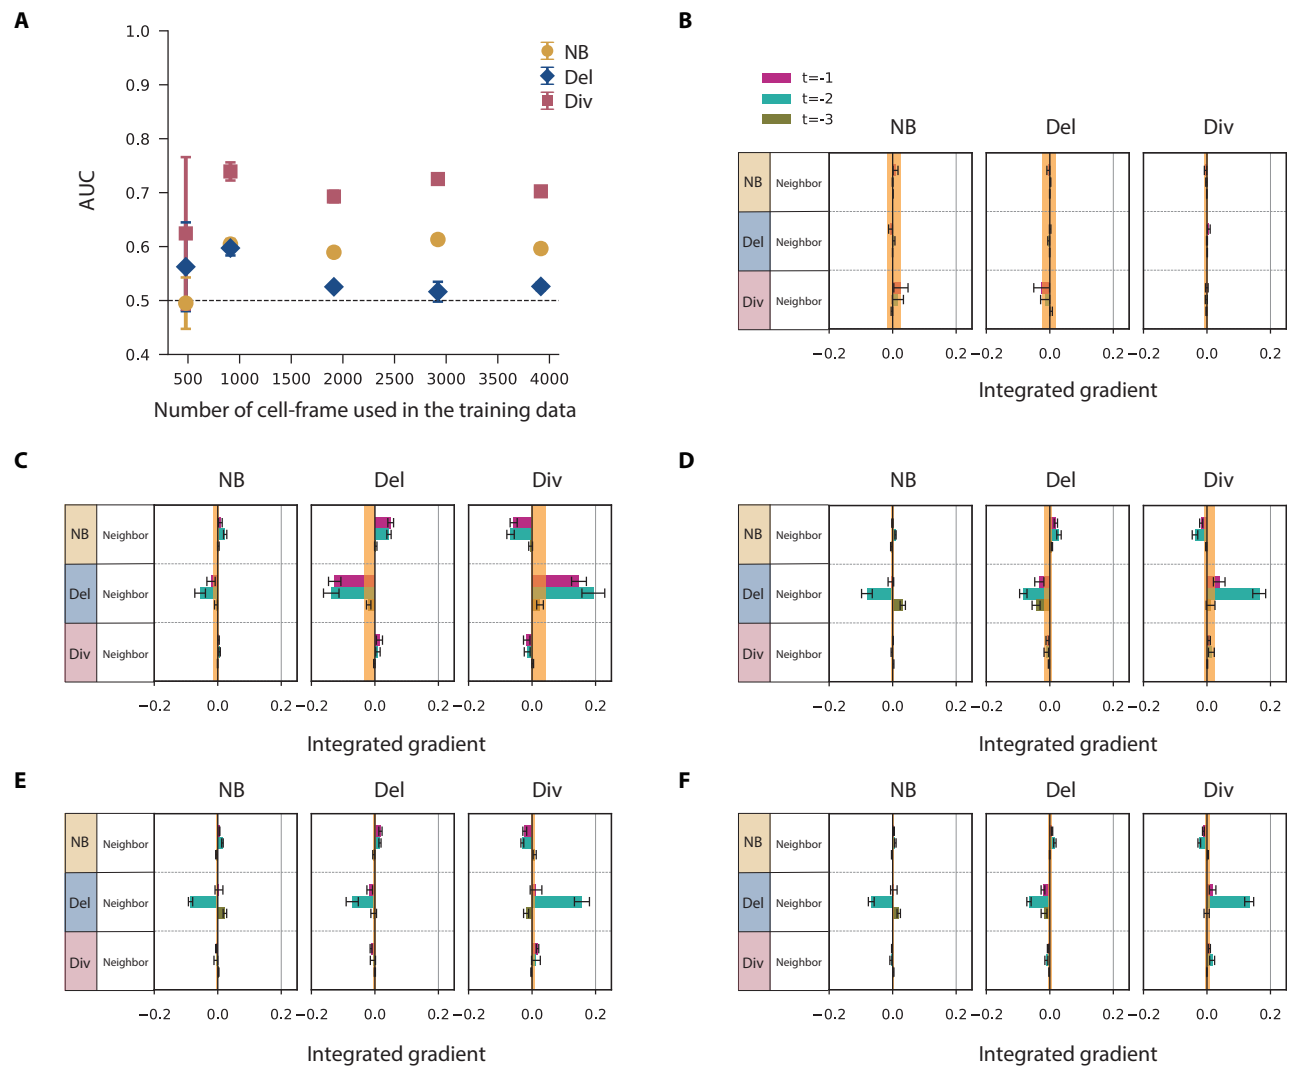


Fig. S11. Data size dependency on the prediction and attribution. The cell external model with mean aggregation is applied to the simulation data from the delamination-induced division setup with the feature condition: (NFB, Random)=(+, +). **(A)** The AUC for the different number of cell-frames in the training data obtained by averaging the AUC over six trained models. Error bar: standard deviation. **(B-F)** The attribution is shown for the different numbers of cells: **(B)** 479, **(C)** 911, **(D)** 1913, **(E)** 2920 and **(F)** 3917. The result in the main text is for 2920 cells in the training data. The IG averaged over six trained models is shown for each pooled feature. Error bar: standard error. The upper and lower values of the IG of the random feature are shown as the orange zone.

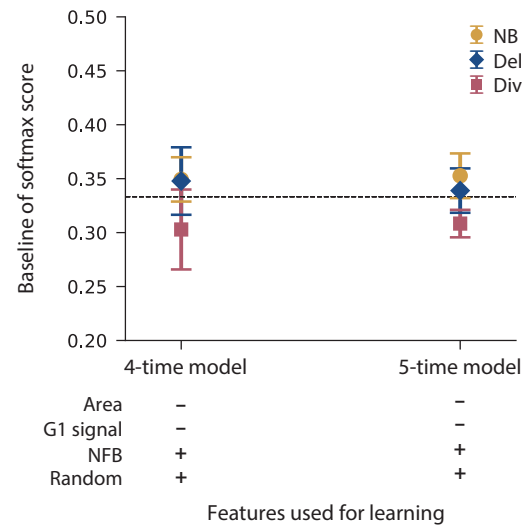


Fig. S12. **Baseline softmax score.** The baseline softmax score, which is the softmax score for null-graphs, is shown for the four and five-time cell external model with mean aggregation for the hind paw data. The feature condition is (Area, G1 signal, NFB, Random)=(-, -, +, +). The horizontal dashed line indicates the target baseline softmax score 1/3.

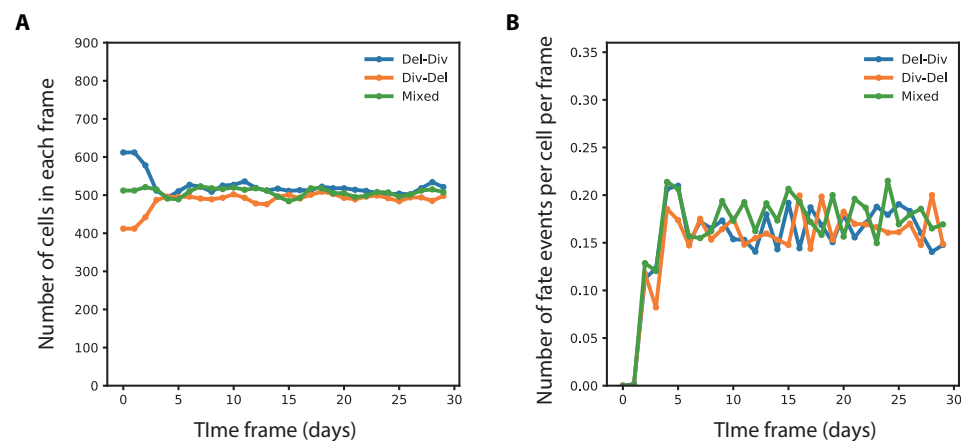


Fig. S13. **Numerical simulations of the homeostatic tissue model.** (A) Time-evolution of the number of cells and (B) the number of fate events in the simulations of the Del-induced Div model, Div-induced Del model, and the mixed model. We used the data from 15 to 30 days in these simulations for the GNN analyses.