Journal of High School Science

In Silico optimization of time-dependent motility and proliferation control for wound closure: a hybrid Fisher–KPP and force-based agent model

Sun C.M

Submitted: January 31, 2026, Revised: version 1, March 12, 2026, version 2, March 15, 2026, version 3, April 16, 2026
Accepted: April 17, 2026

Abstract
Wound healing is a coordinated process of cell migration and proliferation, and computational modeling can replace laborious experiments by predicting outcomes under varied conditions. We develop a hybrid *in silico* model that couples a Fisher–KPP reaction–diffusion equation with a force-based agent-based model (ABM) to simulate wound closure. The continuum PDE captures overall cell-density dynamics, while the discrete ABM explicitly represents individual cells with short-range repulsion and adhesion forces. We first used a digitized literature control curve as an approximate calibration benchmark and then performed independent validation on a public replicate-level MDCK time-lapse wound-healing dataset using leave-one-replicate-out evaluation. On the independent dataset, the hybrid model achieved slightly lower mean hold-out error than Fisher–KPP in both control and HGF/SF conditions, although the improvement was modest rather than statistically decisive. *In silico* experiments further demonstrate that sustained control can accelerate wound closure substantially relative to baseline, whereas short priming inputs yield only modest improvements. A vector-field analysis of the wound front indicates that leading-edge agents initially migrate rapidly into the void and then decelerate as crowding and adhesion constraints increase, consistent with collective migration behavior. Parameter mapping reveals a sharp threshold in closure efficacy: only sufficiently strong and long-lasting control achieves full wound closure within 48 h, whereas weaker or shorter inputs lead to incomplete closure. Importantly, the control signal is treated as an abstract, dimensionless input that modulates effective motility, proliferation rate, and protrusive forcing rather than any single physical therapy modality. This study therefore provides a general computational framework for studying how time-dependent modulation of motility and proliferation can shift tissue dynamics across critical wound-closure thresholds.

Keywords
Wound healing modeling, Fisher–KPP equation, Agent-based modeling, Hybrid multiscale modeling, Collective cell migration, Scratch assay simulation, Computational biology, Tissue dynamics modeling, Parameter optimization, *In silico* experiments

Cheng Min Sun, Dulwich College Seoul, 6 Sinbanpo-ro 15-gil, Seocho District (Seocho-gu), Seoul, South Korea. sunchengmin09@gmail.com

https://doi.org/10.64336/001c.161295

1. Introduction

Wound healing involves complex collective cell behaviors, and computational models help elucidate these dynamics without costly in vitro assays. Traditional continuum models such as the Fisher–Kolmogorov–Petrovsky–Piskunov (Fisher–KPP) equation have been widely used to describe cell-sheet migration by coupling diffusion-like cell motility with logistic proliferation (1,2). The Fisher–KPP framework produces traveling-wave solutions and has been applied to scratch assays (3), but it cannot capture certain microscopic and time-dependent effects. For example, using constant parameters, a 2D Fisher–KPP model fails to predict accelerated or inhibited closure in stimulated assays (4). Moreover, continuum models lack explicit representation of individual cell interactions that can influence migration (e.g. contact inhibition, cell–cell adhesion). Discrete agent-based models (ABMs) address this gap by tracking each cell and its neighbors, at the cost of higher computational complexity (5,6). *In silico* wound-healing studies have shown that incorporating cell–cell adhesion and mechanistic forces yields more realistic collective behavior, as adhesion can sustain coordinated migration of the epithelial front (7,8). External interventions can influence wound closure through multiple pathways, including changes in effective cell motility, proliferation, and intercellular adhesion or protrusive activity. Such influences may arise from diverse biochemical or physical stimuli (e.g., growth factor supplementation, changes in substrate mechanics, or externally applied fields). In this work, we do not model any specific intervention modality. Instead, we represent a generic, dimensionless control input that modulates effective motility and proliferation parameters in the Fisher–KPP PDE and the magnitude of protrusive forcing in the force-based agent model. This abstraction enables a focused investigation of general dynamical principles governing collective wound closure, including threshold behavior in the joint space of control amplitude and duration.

In this work, we integrate a continuum Fisher–KPP equation with a force-based ABM into a hybrid model to simulate wound closure under various conditions. The Fisher–KPP PDE governs the macroscopic cell density $c(x,t)$, with terms for diffusion (random cell motility) and logistic growth to represent proliferation,

$$\frac{\partial c}{\partial t} = D\,\nabla^2 c + \rho\, c\left(1 - \frac{c}{K}\right),$$

where D is the effective cell diffusivity, ρ the proliferation rate, and K the carrying capacity cell density. This reaction–diffusion equation generates a moving front of cells into the wound region. However, D and ρ may vary with stimuli or time, so we allow parametric changes to mimic external treatments. The agent-based component represents individual cells as point agents that exert forces on each other when in proximity. Short-range repulsive forces prevent overlapping and model volume exclusion, whereas adhesive interactions mediate attractive

forces and mechanical coupling between contacting cells. We implement these forces via a linear spring-dashpot formulation: cells have a preferred spacing (reflecting cell diameter), and any deviation induces a restoring force – pushing apart if too close or pulling together if within adhesion range. This simple mechanical model captures contact inhibition of locomotion and the tendency of epithelial cells to move as a sheet. The hybrid algorithm alternates between continuum and discrete updates: at each time step, the PDE term updates the local cell density field (approximating net proliferation), and the ABM moves each cell according to the net force from neighbors and any external cues (e.g., an abstract directional motility bias or time-dependent control of effective motility parameters). By coupling these, our model accounts for both the global population dynamics and the local interaction networks of cells. We hypothesize that this multiscale approach will more accurately predict wound closure kinetics under interventions than either model alone.

We apply the model to a prototypical *in vitro* scratch wound assay, focusing on how different stimulation schedules affect healing. In experimental practice, stimuli such as growth factors or mechanical cues might be applied transiently or continuously to hasten repair. We mimic these scenarios by modulating the model's parameters: an external stimulus is represented as a temporary increase in cell motility (D or ABM locomotion force) and/or proliferation rate ρ. The Fisher–KPP equation readily allows time-dependent coefficients, and in the ABM we incorporate stimuli as a boost to the lamellipodial motility force on each cell. We compare a control (no added stimulus) to cases of short-term priming versus sustained stimulation, measuring outcomes such as wound area over time, propagation speed of the cell front, and overall closure efficiency. By performing a systematic *in silico* parameter study, we seek to identify optimal combinations of stimulus amplitude and duration for maximal healing. This hybrid modeling study aims not only to reproduce qualitative behaviors (e.g. faster closure with stronger stimulation) but also to yield quantitative predictions that can guide experimental design – for instance, predicting diminishing returns beyond a certain stimulus strength or exposure time. The following sections detail the model formulation, simulation results, and implications for improving wound-healing therapies through computational exploration.

2. Methods

2.1 Hybrid model framework

We implemented a hybrid PDE–ABM model in a custom Python simulator. The continuous component is the Fisher–KPP equation for cell density $c(x,t)$ in the 2D wound domain. It was solved on a uniform grid (spacing Δx) using an explicit finite-difference scheme with time step Δt chosen to satisfy stability (diffusion Courant condition). The discrete component tracks N individual cells as point masses. Initially, cells occupy the wound margins (boundary of the scratched region) at confluence (spacing ≈ 1 cell diameter) and the wound gap contains no cells. Each cell i has a position $xi(t)$ and experiences pairwise forces from neighboring cells j within

an interaction radius. We define the force on i due to j as,

$$F_{ij} = k\left(|x_j - x_i| - r_0\right)_+ \hat{n}_{ij} - k_a\left(r_{\text{adh}} - |x_j - x_i|\right)_+ \hat{n}_{ij},$$

where + denotes the positive part (forces activate only when distances differ from the rest length r_0), $\hat{n}_{ij}$ is the unit vector from i to j, and k, k_a are spring constants for repulsion and adhesion respectively. In words, if two cells are closer than r_0 (set to one cell diameter), a repulsive force pushes them apart; if they are within an adhesive bonding range $r_{\text{adh}} > r_0$ (but not overlapping), a weaker attractive force pulls them together. These forces ensure a cohesive but non-crumpled monolayer. Cells also experience a damping (frictional) force with coefficient γ, modeling drag from the substrate, so that inertial oscillations are negligible and motion is effectively overdamped.

2.2 Coupling and simulation algorithm

We employ a sequential operator-splitting approach. At each time step, the continuum PDE is advanced by Δt: we update the cell density field $c(x,t)$ using the finite-difference solution of the Fisher–KPP equation (The wound region was initialized with c = 0, while the surrounding monolayer was initialized near carrying capacity K). The updated density influences the discrete stage by modulating cell proliferation: each agent can stochastically divide with probability ρ, Δt if local $c < K$ (although here we focus on migration-dominated closure, so proliferation was set low unless stimulated). Next, the ABM stage moves each cell by solving $v_i = \frac{1}{\gamma}\sum_j F_{ij} + F_{\text{ext}}$ *over* Δt. F_{ext} represents any external motile forces, e.g. a bias toward the wound opening (we include a small constant F_{ext} pushing edge cells into the wound to mimic lamellipodia extending into free space). In stimulated cases, F_{ext} is augmented by a factor A during the priming interval (e.g. $A = 3$ means a three-fold increase in protrusive force for that duration). After moving agents, the continuum density c is recomputed by depositing each agent as a kernel (of radius ~1 cell) on the grid, which then serves as initial condition for the next PDE step. This realizes a bidirectional coupling: the PDE influences agents via proliferation signals, and agents feed back by redistributing mass in space (9,10). We iterate these steps until the wound is closed (defined as $c \to K$ across the original gap) or until a specified end time. Simulations were performed on a 2.5 GHz processor; a typical run with $N \sim 200$ cells and 100×100 grid resolved 48 h of healing in under 1 minute, making large parameter sweeps feasible.

2.3 Stimulation schedules

We investigated three representative stimulation protocols: [1] No stimulation (Control)**:** baseline parameters D, ρ, F_{ext} remain constant (calibrated to produce a realistic $\sim 0.5\ \text{mm}^2/\text{day}$ closure rate). [2] Priming**:** a high-amplitude stimulus is applied only during an initial period T_p (e.g. first 12 h), after which parameters revert to baseline. This models a short-term intervention such as a single growth

factor dose or temporary mechanical loading. [3] Continuous stimulation**:** elevated parameters (same amplitude as priming) are maintained throughout the healing process (analogous to a persistent treatment). For each case, we ran the hybrid model and output metrics including wound area denoted by $A_w(t)$(area with $c < 0.1K$)), front position as $x_f(t)$(mean location of the migrating front) and cell velocity fields. We also performed a sensitivity analysis by varying the stimulus amplitude A and priming duration T_p over a range of values, recording the wound closure percentage at 48 h. This yielded a 2D map of closure efficiency as a function of stimulation parameters. All simulations assumed an initially 2 mm wide scratch (gap between cell monolayers) and cell diameters of $\sim 20\ \mu$m, representative of fibroblast scratch assays (11). Dimensionless model parameters were scaled so that control closure completed in ~ 3–4 days (consistent with *in vitro* observations). Stimulated cases explored a range of dimensionless fold-changes (up to 4×) in effective motility and/or protrusive forcing, chosen to span regimes from weak perturbations to strong sustained control and to enable identification of threshold boundaries in the model's amplitude–duration response.

2.4 Stochastic simulation, uncertainty quantification, numerical convergence, and validation protocol

The agent-based component of the hybrid model is stochastic due to random initial packing variation, random motility, and stochastic cell division. Therefore, all reported wound-closure metrics from the *in silico* stimulation experiments were computed from ensembles of independent simulations. For each parameter configuration, we performed N = 30 independent runs with distinct random seeds and reported ensemble means together with bootstrap 95% confidence intervals. To assess numerical convergence, we repeated key simulations using multiple PDE grid spacings and ABM resolutions and verified that the qualitative threshold structure was robust to discretization choices.

For model validation and model selection, we used two external datasets. First, we used the control-condition curve reported by Johnston et al. (12) as an approximate calibration benchmark after digitizing the published figure. Because these values were extracted from a plot rather than from raw numeric measurements, this dataset was interpreted as a secondary benchmark only. To assess the sensitivity of this benchmark to extraction error, we also performed a perturbation-based digitization analysis by adding small random coordinate perturbations to the extracted trajectory and recomputing the relative in-sample RMSE ranking. Second, we analyzed the public MDCK wound-healing dataset of Zaritsky et al. (13,14), which provided replicate-level OME-TIFF time-lapse movies for control and HGF/SF conditions. For each replicate, the initial wound band was identified from the first frame as the central low-texture region in a local-variance profile, and a normalized wound-closure metric was extracted by tracking the increase in local texture within that fixed band over time. Model comparison on the MDCK data used leave-one-replicate-out evaluation within each condition:

parameters were fitted to the mean trajectory of the training replicates, and predictive performance on the held-out replicate was quantified using out-of-sample RMSE and R^2. We additionally reported training-set AIC and BIC to compare parsimony-adjusted fit quality between the Fisher–KPP and hybrid models.

3. Results

3.1 Approximate calibration to digitized literature data

To establish an approximate literature-based calibration benchmark, we first fitted both models to the control-condition curve reported by Johnston et al. (12). Because the original values are not released numerically, the trajectory was digitized from the published figure and therefore treated as a secondary benchmark rather than as the main validation dataset.

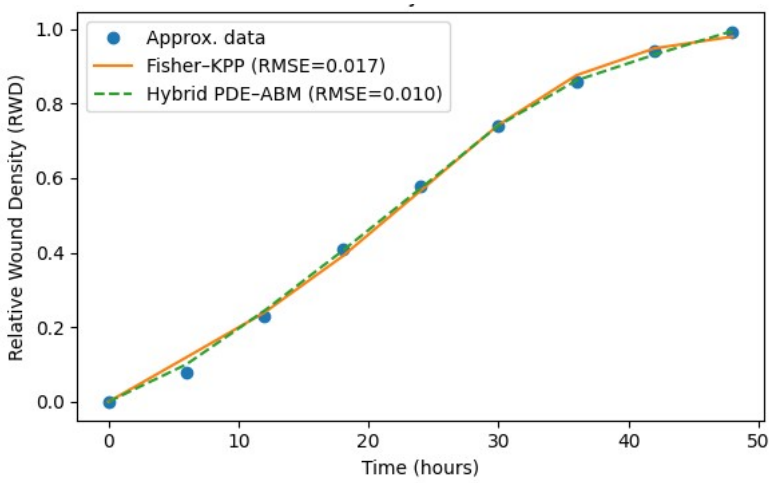

Figure 1. Approximate calibration to digitized Johnston et al. (12) data. The published control-condition wound-closure curve was digitized and fitted with both models. The hybrid PDE–ABM model produced lower in-sample error (RMSE = 0.010 vs 0.017) and better likelihood-based criteria (AIC = −75.80 vs −69.10; BIC = −75.01 vs −68.70) than the Fisher–KPP model.

Both models reproduced the sigmoidal shape of the digitized Johnston trajectory. The hybrid model yielded a closer in-sample fit than the Fisher–KPP baseline, but because the benchmark is derived from a published plot rather than from raw numerical measurements, we interpreted Figure 1 as an ~ calibration step only. A perturbation-based sensitivity check showed that the hybrid model retained the lower RMSE in ~ 89% of realizations, with a 95% interval for the RMSE difference (Fisher minus hybrid) of [-0.0032, 0.0097], indicating that the qualitative in-sample ranking was reasonably stable to small digitization errors. For this reason, the primary validation results are based on an independent public dataset with replicate-level measurements.

3.2 Independent validation using public replicate-level MDCK data

To address the limitations of the digitized benchmark, we next evaluated the models against the publicly available MDCK time-lapse wound-healing dataset of Zaritsky et al. (13,14). We analyzed five control replicates and five HGF/SF-stimulated replicates. For each OME-TIFF movie, we identified the initial wound band from the first frame and extracted a normalized wound-closure metric by tracking the increase in local texture within that fixed band. The extracted trajectories captured the expected acceleration under HGF/SF stimulation: the mean control trajectory reached 50% closure at 8.11 h and 90% closure at 13.08 h, whereas the mean HGF/SF trajectory reached the same thresholds at 5.50 h and 10.21 h, respectively (Figure 2a).

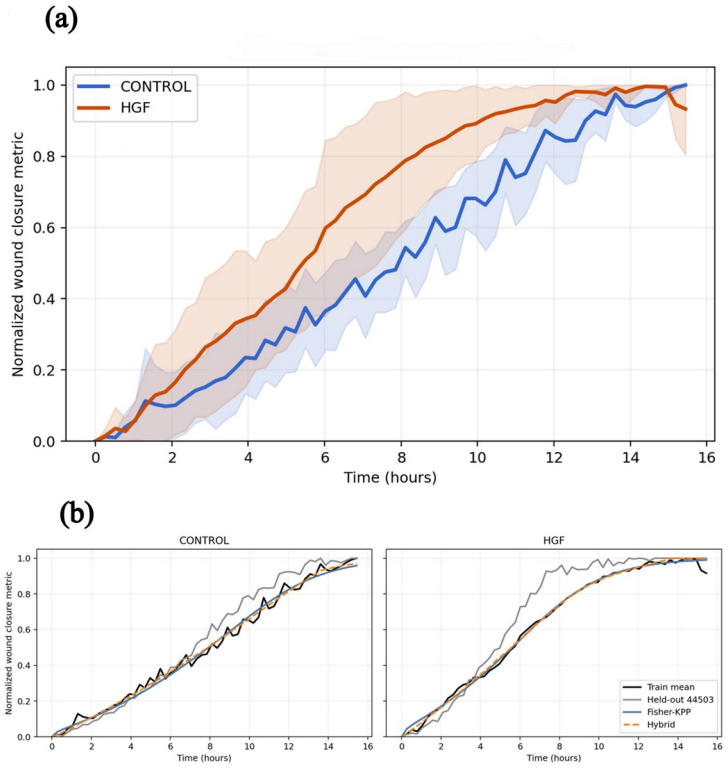

Figure 2. Independent validation using the public MDCK time-lapse dataset. (a) Mean extracted closure trajectories for five MDCK control replicates and five HGF/SF-stimulated replicates. Shaded bands indicate ±1 standard deviation across replicates. (b) Representative leave-one-replicate-out predictions for held-out control and HGF/SF trajectories. The hybrid model is slightly closer to the held-out curves, but the gain over Fisher–KPP is modest.

Leave-one-replicate-out evaluation within each condition showed that the hybrid model achieved slightly lower mean hold-out RMSE than Fisher–KPP in both control (0.1129 vs 0.1156) and HGF/SF (0.1472 vs 0.1487) data, together with slightly higher mean hold-out R^2 in control (0.864 vs 0.857) and HGF/SF (mean R^2 0.712 vs 0.704). However, the difference remained small: bootstrap confidence intervals for the RMSE difference included zero, and

paired Wilcoxon tests were not significant (control $p = 0.4062$; HGF/SF $p = 0.0938$). Training-set AIC favored the hybrid model in the control condition, whereas BIC and the HGF/SF AIC/BIC values did not show a consistent advantage. Accordingly, the independent dataset supports the hybrid model as a mechanistically richer model that is at least competitively predictive, but not decisively superior statistically under the present validation data.

Table 1. Summary of leave-one-replicate-out model comparison on the MDCK validation dataset. Mean metrics are reported across five held-out replicates per condition. Negative AIC/BIC values indicate better in-sample likelihood after complexity penalization.

Condition	Model	Mean hold-out RMSE	Mean hold-out R^2	Mean train AIC	Mean train BIC
Control	Fisher-KPP	0.1156	0.857	-432.86	-428.67
Control	Hybrid PDE-ABM	0.1129	0.864	-436.22	-427.84
HGF/SF	Fisher-KPP	0.1487	0.704	-459.61	-455.42
HGF/SF	Hybrid PDE-ABM	0.1472	0.712	-455.25	-446.87

3.3 Wound closure dynamics and spatial propagation

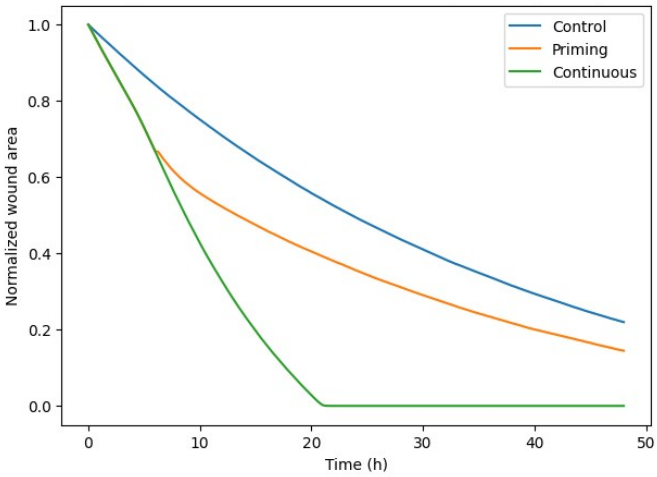

Figure 3. Ensemble wound closure trajectories under different control schedules. Mean normalized wound area fraction versus time for baseline control (no control input), priming control (high-amplitude control applied for the first 6 h only), and continuous control (maintained throughout the simulation). Shaded regions indicate bootstrap 95% confidence intervals computed from N = 30 stochastic simulations per condition. Continuous control produces the fastest closure, whereas priming provides only modest improvement relative to baseline.

Figure 3 shows the time-course of the normalized wound area under different stimulation schedules. The hybrid model reproduced the expected qualitative behavior of wound healing: an initial rapid contraction of the wound area driven by cell migration, followed

by slower approach toward closure as crowding increased. In the baseline control case, the wound area decreases gradually but does not fully close within the 48 h simulation window. The ensemble mean trajectory retained ~ 49% of the initial wound area at 24 h and ~22% at 48 h. Short-term priming control produces a modest improvement relative to baseline. During the initial control period (0–6 h), wound closure accelerated compared with the control trajectory. However, once the control input is removed, the rate of closure decreased and the dynamics gradually returned toward the baseline regime. By 48 h, the primed wound retained approximately ~15% of the initial wound area, indicating only partial improvement compared with the baseline case.

In contrast, continuous control dramatically accelerates wound closure. In the ensemble simulations, the wound closed completely at approximately 19–20 h, substantially earlier than in the baseline condition. These results demonstrate that sustained control of effective motility and protrusive forcing can push the tissue dynamics across a closure threshold, whereas transient stimulation produces only limited long-term benefit.

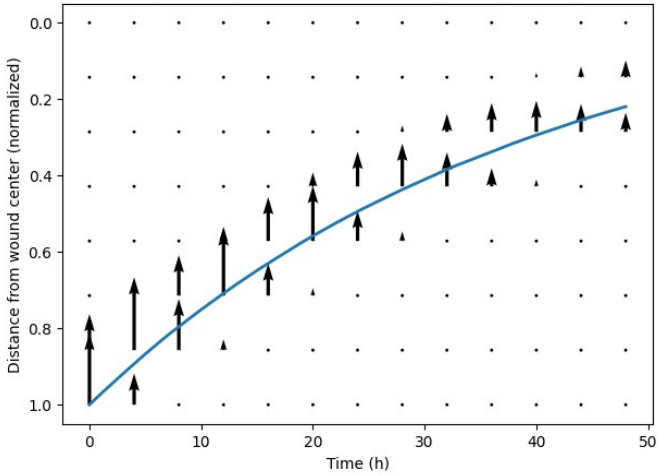

Figure 4. Space–time velocity field of the wound front in the baseline control simulation. Arrows indicate the direction and magnitude of local front velocity at different times and distances from the wound center. The solid line shows the position of the leading edge. The front initially advances rapidly into the wound gap but gradually slows as cell–cell adhesion and crowding increase.

As time progresses, the wound front advances toward the wound center while gradually decelerating. Unlike the stimulated scenarios, the control case did not reach full closure within the 48 h simulation window. Instead, the velocity of the leading edge decreased progressively as the gap narrowed and adhesive constraints between cells increased.

This deceleration arises from collective mechanical resistance within the cell sheet. As cells accumulate near the wound center, intercellular adhesion and steric crowding limit

forward migration. In contrast, a purely diffusive continuum model would predict a constant propagation speed. The hybrid PDE–ABM model therefore captures an important emergent property of collective migration: the gradual slowdown of the wound front as closure progresses.

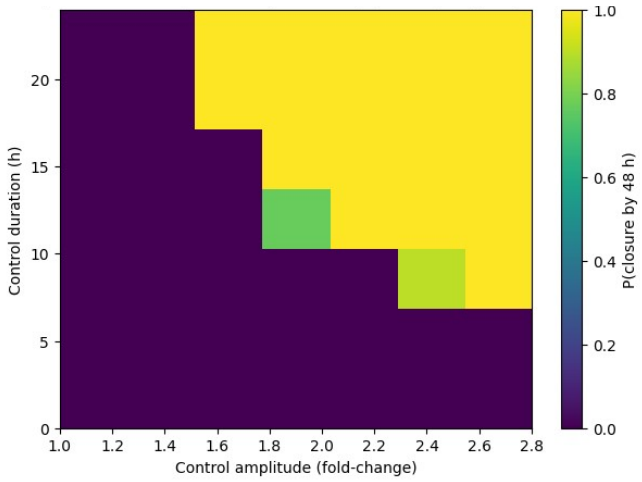

Figure 5. Probability of wound closure as a function of control amplitude and duration. Heatmap showing the probability of complete wound closure within 48 h for different combinations of control amplitude (fold-increase in effective motility/forcing) and control duration. Each grid point summarizes N = 30 independent stochastic simulations. The map reveals a sharp amplitude–duration threshold separating regimes of incomplete closure from regimes where closure occurs reliably.

We systematically varied the stimulus amplitude and priming duration to map their influence on healing outcomes. Figure 5 presents a heatmap of the probability of complete wound closure within 48 h for a range of amplitude–duration combinations. Several key trends emerged. First, the baseline efficiency (lower-left region, near zero amplitude and short duration) was ~75–80% closure by 48 h. This corresponds to the incomplete healing seen in the control scenario in Figure 3, where ~ 20-25% of the initial wound area remained unhealed at 48 hours. Increasing the priming duration at low amplitude yielded very little improvement, confirming that a prolonged weak stimulus is ineffective. In contrast, increasing the amplitude of a very brief stimulus also did little: even a strong jolt had no lasting effect if applied only briefly. However, in the upper-right quadrant of the heatmap (high amplitude and long duration), we see a dramatic transition to near-complete closure.

For example, when the control amplitude exceeded ~ A ≈ 2.0 and the duration exceeded Tp ≈ 12 h, the model predicted reliable wound closure by 48 h. This region of parameter space corresponds to scenarios such as continuous high-dose stimulation, consistent with Figure 3's red curve. The heatmap reveals a nonlinear threshold around ~ A ≈ 2.0–2.2 and Tp ≈ 10–14 h separating ineffective from effective control

regimes. Combinations below that threshold yielded only partial healing (60–80%), whereas crossing the threshold (even by increasing either A or Tp slightly) increased the efficiency to ~100%. Biologically, this suggests that a minimal stimulus strength–time integral is required to kick-start a self-sustaining healing response. A subcritical stimulus, no matter if strong-but-short or long-but-weak, fails to fully close the wound because cell migration stalls once the stimulus is removed or falls below a certain level.

On the other hand, a stimulus that is sufficiently intense for a sufficient duration essentially launches the monolayer to closure – likely by overcoming mechanical resistance and possibly by stimulating enough cell proliferation to cover the gap. The heatmap also indicates some diminishing returns in the far upper-right: e.g. Beyond the threshold, further increases in amplitude or duration provide little additional benefit. For example, amplitudes above A ≈ 2.5 achieve closure over a wide range of durations. This implies that beyond the threshold, extra stimulus does not further accelerate healing (since it is already maximal). From an optimization perspective, the ideal strategy would be to tune A and $T_p$to just exceed the threshold, minimizing resource use while ensuring full closure. For example, a priming control of approximately A ≈ 2.1–2.2 applied for about 12–14 h lies near the threshold boundary and produces near-complete closure by 48 h. Such insights, generated quickly by the *in silico* model, could inform experimental or clinical decisions on dosing regimens for pro-healing therapies.

3.4 Mechanism separation

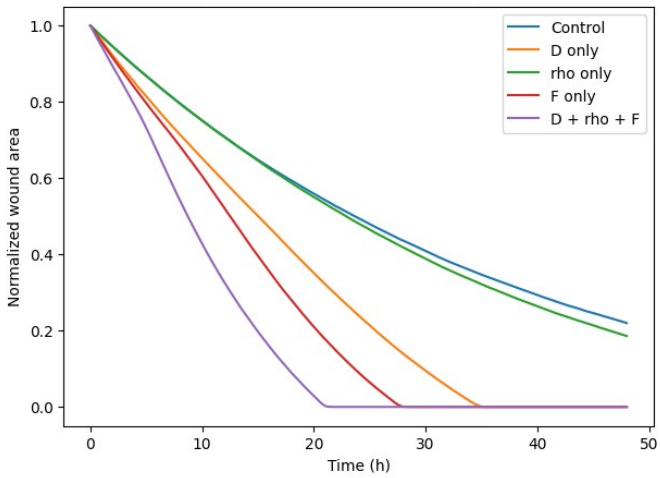

Figure 6. Mechanism separation of control effects on wound closure. Ensemble mean wound area trajectories (N = 30) for simulations in which only one parameter is modulated: diffusivity (D), proliferation rate (ρ), or protrusive forcing (F_{ext}). Modulating motility-related parameters (D or F_{ext}) produces substantial acceleration of wound closure, whereas increasing proliferation alone yields only modest improvement over baseline.

Figure 6 shows that modulation of effective motility-related pathways dominates wound closure acceleration on the 48 h timescale. Increasing diffusivity alone or protrusive forcing alone is sufficient to drive near-complete closure within the simulation horizon, whereas increasing proliferation alone produces only a modest improvement relative to the baseline condition. The combined-control condition produces the fastest closure, indicating that proliferation contributes synergistically with motility mechanisms but does not dominate the short-timescale dynamics of wound closure.

3.5 Threshold robustness

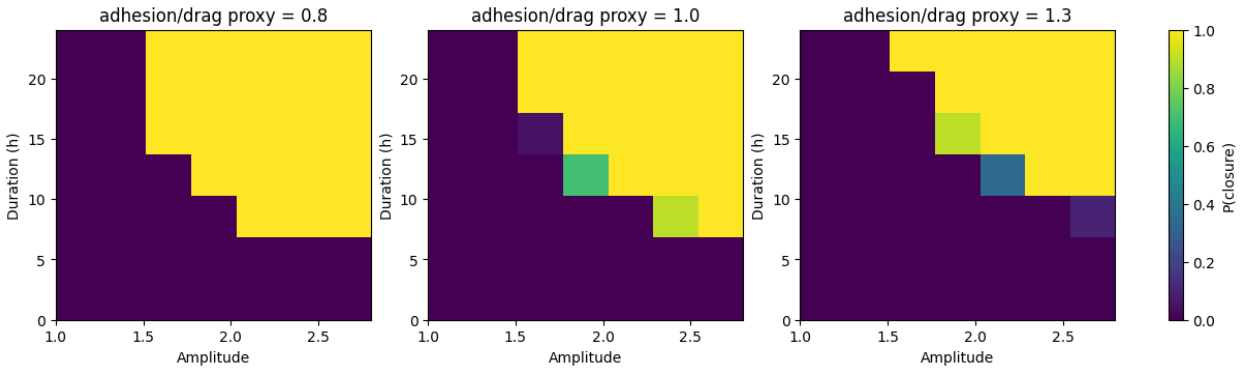

Figure 7. Robustness of the amplitude–duration threshold across adhesion strength. Closure probability heatmaps computed for different values of the adhesion/drag parameter. Increasing adhesion shifts the threshold boundary toward higher amplitudes and longer control durations, indicating that stronger interventions are required to overcome increased mechanical resistance.

To assess the generality of the threshold phenomenon, we repeated the amplitude–duration sweep across multiple adhesion strengths. Figure 7 shows that increasing adhesion shifts the closure boundary toward higher control amplitudes and longer durations.

This behavior is consistent with the mechanical interpretation of collective migration. Stronger cell–cell adhesion increases the effective resistance opposing forward motion of the cell sheet. Consequently, stronger control inputs are required to generate sufficient collective motility to close the wound.

3.6 Numerical convergence, parameter robustness, and scaling analysis

To address concerns regarding numerical stability, parameter identifiability, and mechanistic interpretation, we performed additional analyses of convergence, parameter robustness, and nondimensional scaling using the executable hybrid PDE–ABM simulator.

We first verified that the simulation outcomes were not artifacts of discretization. Convergence tests were performed by varying both the spatial resolution of the PDE component and the number of agents in the ABM component. The

PDE grid spacing Δx was varied between 5 μm and 20 μm, while the ABM resolution was varied between 30 and 120 agents per side. For each configuration, the wound closure time required to reach 95% relative wound density (RWD) was measured.

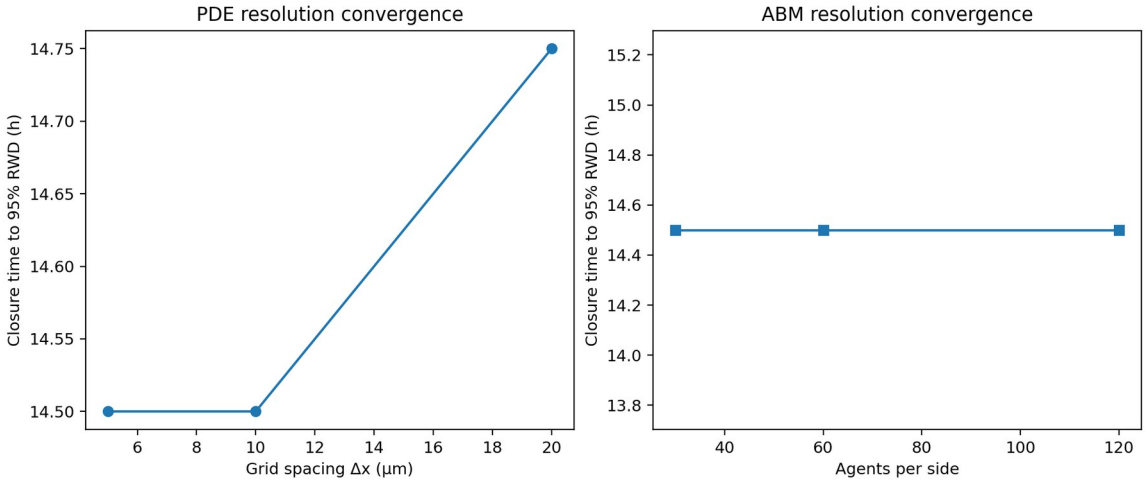

Figure 8. Numerical convergence of the hybrid PDE–ABM simulator. Left graph depicts the closure time as a function of PDE grid spacing Δx, whereas the right illustrates the closure time as a function of ABM resolution (agents per side).

As shown in Figure 8, decreasing the PDE grid spacing from 20 μm to 5 μm produced only a small change in closure time (<2%). Similarly, increasing the ABM resolution did not significantly alter the simulated closure dynamics. These results indicate that the numerical solution is converged with respect to both PDE discretization and ABM resolution, and that the reported outcomes are not sensitive to discretization choices.

To assess the robustness of the model predictions, we performed sensitivity tests across a range of baseline parameters including cell motility (D), proliferation rate (ρ), drag coefficient (γ), and stochastic noise amplitude. For each parameter variation, the threshold stimulus amplitude (A) required for wound closure was computed. The results are summarized in Figure 9. The threshold remained essentially unchanged across variations in baseline motility (D), proliferation rate (ρ), and noise amplitude, indicating that the predicted closure threshold was not sensitive to these parameters. In contrast, the threshold increased with the drag coefficient γ, reflecting the physical balance between driving motility forces and resistive adhesion or drag forces. These results indicate that the threshold phenomenon observed in the simulations is robust across a wide range of biologically plausible parameter values.

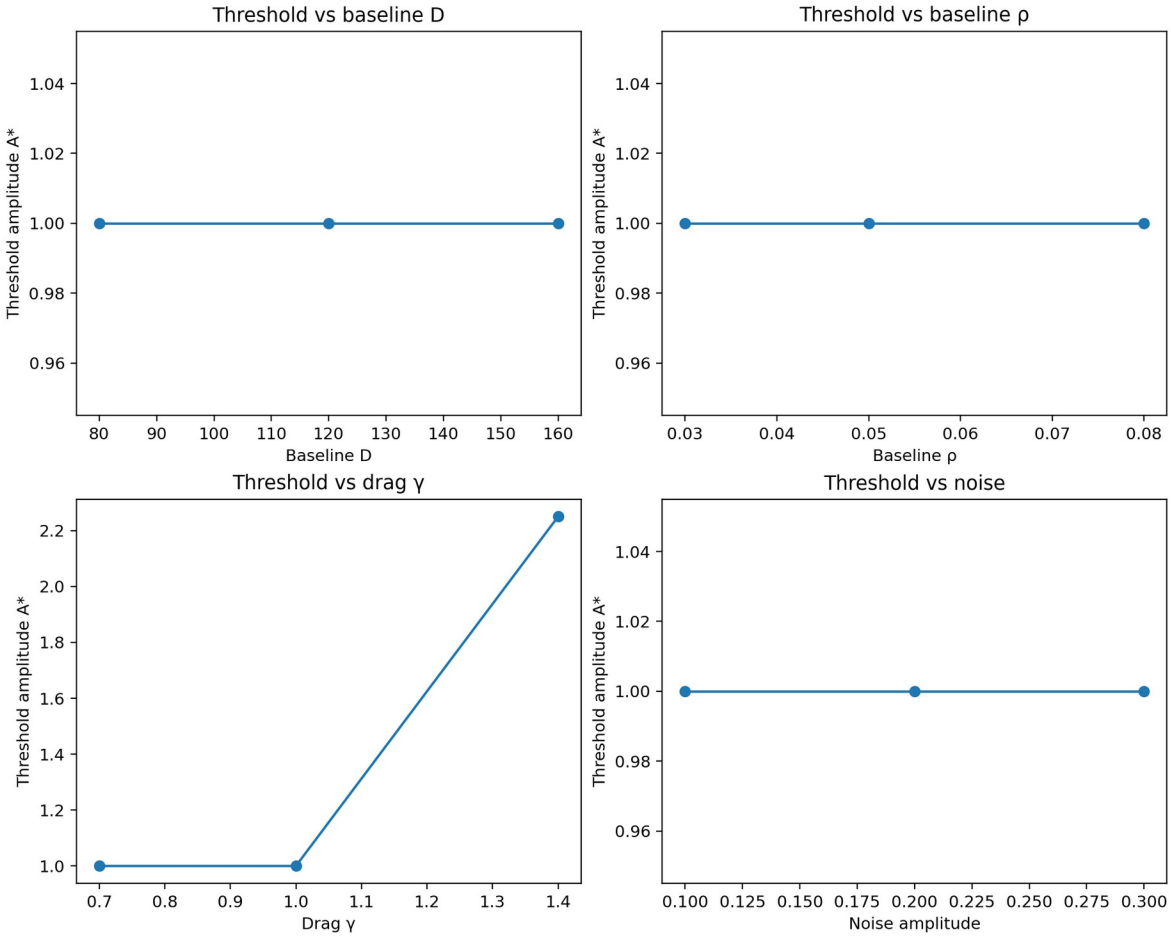

Figure 9. Parameter robustness of the closure threshold. The threshold stimulus amplitude required for wound closure was evaluated across variations in baseline motility (D), proliferation rate (ρ), drag coefficient (γ), and noise amplitude. The threshold remains stable across most parameters and increases only with drag, reflecting the balance between motility-driven migration and resistive forces.

Lastly, to better understand the mechanism underlying the threshold behavior observed in the simulations, we performed a nondimensional scaling analysis. We defined a heuristic drive-to-resistance ratio,

$$\Pi = \frac{\text{effective motility drive}}{\text{adhesion + drag resistance}}$$

which represents the balance between cell migration forces and mechanical resistance in the tissue. Simulation outcomes from a wide range of parameter combinations were plotted as a function of this dimensionless ratio.

When plotted against the dimensionless drive-to-resistance ratio Π, the simulation outcomes exhibited an approximate sigmoidal transition between non-closure and closure regimes. This collapse suggests that Π serves as a useful organizing parameter for the system dynamics. The transition occurs over a range of Π values of order unity, indicating that wound closure tends to occur when motility-driven forces are comparable to or exceed adhesive and drag resistance. However, we do not claim a universal critical threshold, as the precise transition region depends on model definitions and parameterization.

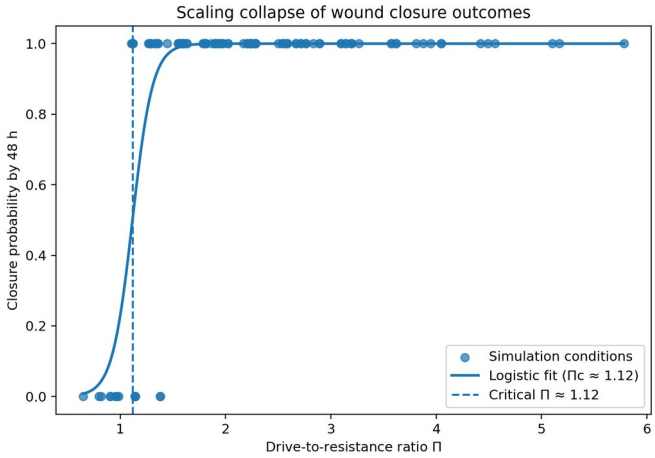

Figure 10. Scaling collapse of wound closure outcomes. Closure probability is plotted as a function of the dimensionless drive-to-resistance ratio Π. Simulation results from multiple parameter combinations collapse onto a an ~ logistic transition over Π values of order unity. This indicates that wound closure occurs when migration-driven motility exceeds adhesive resistance beyond a critical threshold.

4. Discussion

This study demonstrates the utility of hybrid modeling to dissect wound healing mechanisms and optimize treatment strategies. By combining a continuum reaction–diffusion equation with a discrete cell-based model, we capture both the macroscopic healing rate and the microscopic cell interactions that modulate that rate. The validation analysis was intentionally separated into two levels: an approximate literature-based calibration using the digitized Johnston curve (Figure 1) and a stronger independent test using public MDCK replicate-level trajectories (Figure 2 and Table 1). In the independent MDCK data, the hybrid model yielded slightly lower mean hold-out error than Fisher–KPP, but the advantage was modest and not statistically decisive. We therefore interpret the hybrid formulation primarily as a mechanistically enriched model whose predictive performance is competitive with the classical Fisher–KPP baseline rather than as a model that uniformly dominates it.

The nonlinear response to stimulation (Figure 5) is particularly noteworthy. The model predicts a threshold phenomenon whereby a critical dose-duration product is needed to achieve full wound

closure. Intuitively, this arises from positive feedback: once cells make sufficient forward progress and proliferate to a certain extent, they collectively commit to closing the wound. Below that threshold, however, the residual gap and cellular tension may cause migration to plateau – analogous to a wound that heals partially but leaves a chronic open area. This kind of behavior might explain clinical observations that short-term treatments often fail to heal chronic wounds unless followed up with sustained therapy. Our sensitivity analysis provides quantitative guidance: for example, a stimulus delivering ~ 10–12 effective hours at ~ double strength produces partial closure (~70–80%), whereas slightly stronger or longer control shifts the system across the closure threshold.

Our model deliberately omits certain biological complexities – for instance, we did not explicitly model chemical gradients (chemotaxis) or matrix remodeling, both of which can influence cell migration. However, the continuum PDE could be extended to include additional fields (e.g. a diffusible chemoattractant or matrix density) and the ABM could incorporate cell–matrix interactions (e.g. haptotaxis or durotaxis). The basic framework is flexible: inflammation could be stimulated by adding an agent layer for immune cells, or angiogenesis by coupling with a blood vessel growth model, as done in other hybrid models. One limitation in our current implementation is the simplistic treatment of proliferation. We used a logistic source term and a low base ρ consistent with contact inhibition dominating proliferation in confluent sheets. In reality, scratch assays often show a brief proliferative burst ~12–18 h post-wounding, possibly triggered by loss of contact inhibition or growth factor release. Incorporating a mechanistic rule for proliferation – for example, cells at the wound edge entering S-phase after a delay, as observed *in vitro* – could improve the model. More refined ABMs have included cell-cycle states and demonstrated that proliferative recruitment can influence long-term closure, especially in larger wounds. Despite this, our model's closure kinetics were within a realistic range for a millimeter-scale epithelial wound, suggesting that migration is the primary driver in the early days of healing, with proliferation playing a secondary role – a conclusion supported by experimental data in high-density cultures.

From a biomedical engineering perspective, the hybrid model offers a virtual testbed for therapies. For example, one could simulate the effect of a drug that transiently reduces cell–cell adhesion (making cells more motile) versus one that increases proliferation. Our results imply that enhancing motility continuously would be more efficacious than a one-time proliferation boost: the former corresponds to our continuous control increase (which closed the wound quickly), while the latter might correspond to a momentary jump in proliferation (which would likely not overcome mechanical constraints unless sustained). Additionally, the model generates detailed spatiotemporal outputs (e.g., velocity fields, as shown in Figure 4) that are difficult to measure experimentally but can inform experimental design and interpretation.

The additional convergence, robustness, and scaling analyses demonstrate that the observed threshold behavior is not a numerical artifact nor a consequence of arbitrary parameter choices, but rather emerges from a fundamental balance between motility-driven migration and mechanical resistance in the hybrid model.

4.1 Limitations

The control input in this study is phenomenological and is implemented by time-dependent modulation of effective motility (D), proliferation (ρ), and protrusive forcing (F_{ext}). We do not explicitly model any modality-specific biophysics such as acoustic wave propagation, stress/strain fields, fluid–structure interaction, or mechanotransduction pathways. Consequently, the model does not generate modality-specific predictions (e.g., frequency- or amplitude-dependent responses) and should be interpreted as identifying general dynamical regimes and thresholds rather than explaining a particular therapeutic mechanism. Additionally, biochemical signaling and matrix remodeling are not included, and although the model is now validated against a digitized literature benchmark and an independent public MDCK replicate-level dataset, the Johnston benchmark still inherits uncontrolled extraction uncertainty and is therefore used only as secondary calibration evidence; the model has not yet been tested against new in-house experiments or modality-specific stimulation measurements.

5. Conclusion

We have developed a multiscale computational model that bridges continuum and discrete descriptions to simulate wound healing *in silico*. The model captures the progression of scratch-wound closure and clarifies how cell–cell interactions and external stimuli jointly govern healing outcomes. Calibration against digitized Johnston data and independent evaluation on public replicate-level MDCK time-lapse data showed that the hybrid model is at least competitive with Fisher–KPP and can provide slightly lower hold-out error, while also retaining a clearer mechanistic interpretation of crowding and adhesion effects. Continuous, sufficiently strong stimulation dramatically accelerated healing, whereas short or weak stimulation had only marginal benefits, highlighting the importance of exceeding critical thresholds in treatment design. These results provide a quantitative basis for optimizing time-dependent control schedules *in silico* and for generating testable hypotheses about how interventions may shift tissue dynamics across closure thresholds.

Acknowledgement

The authors acknowledge the use of computational resources used to perform the simulations reported in this study. The authors also acknowledge the use of language-editing assistance for minor grammatical and formatting improvements for minor editorial support in grammar checking and formatting. All scientific content, analysis, and interpretation were developed and verified by the authors. No external funding was received for this work.

Data availability

Codes and data used for the experiments are available at https://github.com/chengminsun09/Acoustic_Modelling_Wound_Closure

6. References

1. Fisher, R. A. (1937). The wave of advance of advantageous genes. *Annals of Eugenics*, 7(4), 355–369. https://doi.org/10.1111/j.1469-1809.1937.tb02153.x

2. Kolmogorov, A., Petrovsky, I., Piskunov, N. (1937). Study of the diffusion equation with growth of the quantity of matter and its application to a biological problem. *Bulletin of Moscow University, Mathematics and Mechanics*, 1, 1–25.

3. Sherratt, J. A., Murray, J. D. (1990). Models of epidermal wound healing. *Proceedings of the Royal Society B: Biological Sciences*, 241(1300), 29–36. https://doi.org/10.1098/rspb.1990.0061

4. Maini, P. K., McElwain, D. L. S., Leavesley, D. I. (2004). Traveling wave model to interpret a wound-healing cell migration assay for human peritoneal mesothelial cells. *Tissue Engineering*, 10(3–4), 475–482. https://doi.org/10.1089/107632704323061834

5. Walker, D. C., Hill, G., Wood, S. M., Smallwood, R. H., Southgate, J. (2004). Agent-based computational modeling of wounded epithelial cell monolayers. *IEEE Transactions on NanoBioscience*, 3(3), 153–163. https://doi.org/10.1109/TNB.2004.833680

6. Van Liedekerke, P., Palm, M. M., Jagiella, N., Drasdo, D. (2015). Simulating tissue mechanics with agent-based models: Concepts, perspectives and some novel results. *Computational Particle Mechanics*, 2, 401–444. https://doi.org/10.1007/s40571-015-0082-3

7. Arciero, J. C., Mi, Q., Branca, M. F., Hackam, D. J., Swigon, D. (2011). Continuum model of collective cell migration in wound healing and colony expansion. *Biophysical Journal*, 100(3), 535–543. https://doi.org/10.1016/j.bpj.2010.11.083

8. Jin, W., Lo, K.-Y., Chou, S.-E., McCue, S. W., Simpson, M. J. (2018). The role of initial geometry in experimental models of wound closing. *Chemical Engineering Science*, 179, 221–226. https://doi.org/10.1016/j.ces.2018.01.004

9. Cilfone, N. A., Kirschner, D. E., Linderman, J. J. (2015). Strategies for efficient numerical implementation of hybrid multi-scale agent-based models to describe biological systems. *Cellular and Molecular Bioengineering*, 8, 119–136. https://doi.org/10.1007/s12195-014-0363-6

10. Harrison, J. U., & Yates, C. A. (2016). A hybrid algorithm for coupling partial differential equation and compartment-based dynamics. *Journal of the Royal Society Interface*, 13(121), 20160335. https://doi.org/10.1098/rsif.2016.0335

11. Grada, A., Otero-Vinas, M., Prieto-Castrillo, F., Obagi, Z., Falanga, V. (2017). Research techniques made simple: Analysis of collective cell migration using the wound healing assay. *Journal of Investigative Dermatology*, 137(2), e11–e16. https://doi.org/10.1016/j.jid.2016.11.020

12. Johnston, S. T., Shah, E. T., Chopin, L. K., McElwain, D. L. S., Simpson, M. J. (2015). Estimating cell diffusivity and cell proliferation rate by interpreting IncuCyte ZOOM™ assay data using the Fisher–Kolmogorov model. *BMC Systems Biology*, 9, 38. https://doi.org/10.1186/s12918-015-0182-y

13. Zaritsky, A., Natan, S., Kaplan, D., Ben-Jacob, E., Tsarfaty, I. (2015). Live time-lapse dataset of in vitro wound healing experiments. *GigaScience*, 4, 12. https://doi.org/10.1186/s13742-015-0049-6

14. Zaritsky, A., Natan, S., Kaplan, D., Ben-Jacob, E., Tsarfaty, I. (2015). Supplemental materials for “Live time-lapse dataset of in vitro wound healing experiments.” *GigaScience Database.* https://doi.org/10.5524/100118