THE UNIVERSITY *of* EDINBURGH
School of Physics
and Astronomy

Understanding Waddington's Landscape through Bridging-Induced Phase Separation and Transcriptional Dynamics

MPhys Project Report

Nishwal Gora

Submitted for the 40pt MPhys Project course PHYS11016
March 22, 2026

Abstract

Cellular development is often viewed through Waddington's landscape, where differentiated states act as stable attractors. Here, we use a polymer model of bridging-induced phase separation to study how chromatin organisation reshapes transcriptional behaviour at different stages of cellular development. From simulated 3D chromatin and transcription factor dynamics, we construct binary transcription-unit activity traces and analyse them with a telegraph model. Differentiation redistributes activity away from intermediate switching regimes towards locked ON and OFF states, suppresses burst initiation, and concentrates transcription into a restricted subset of units, a process we term fossilisation. Commitment arises primarily through suppression of OFF→ON transitions, while transcription-factor abundance strengthens memory and canalisation during reversal dynamics.

Supervisor: Professor D. Marenduzzo

Personal statement

The first month of this project was devoted primarily to building the necessary background. I read a substantial part of the literature on Bridging-Induced Phase Separation (BIPS), together with textbooks on polymer physics and DNA such as *"The Theory of Polymer Dynamics"* by Doi and Edwards and *"Understanding DNA"* by Calladine et al. Once I had a reasonable grasp of the theoretical framework, I translated much of the baseline BIPS simulation code, originally provided by my supervisor in Fortran, into Python. This involved not only learning to work with LAMMPS and visualisation tools such as VMD, but also deciding which features of the original implementation were essential to preserve and which could be simplified without losing the core physical mechanisms relevant to the project.

During the second month, I began extending the baseline model by increasing the number of transcription-factor species and transcriptional-unit classes along the chromatin. At the same time, I started to formulate the main conceptual direction of the project: I wanted to understand how differentiation reshapes transcriptional dynamics, and this naturally led me to Waddington's landscape as an organising framework. To support this, I read more widely on the biology of differentiation and senescence and then began constructing simplified "toy" chromatin states representing stem-like, differentiated, and senescent organisation. Even at this stage, much of the work involved deciding what should and should not be built into the model. I had to think carefully about which changes ought to be represented as altered chromatin accessibility, which as changes in the transcription-factor pool, and how much biological realism could be sacrificed in order to keep the model interpretable as a physical system rather than an unmanageably detailed biochemical one.

A major part of the project was devoted to identifying an appropriate dynamical regime for the simulations. I did not treat the model parameters as fixed technical inputs, but as part of the scientific problem itself. The aim was to work in a regime that remained physically interpretable while still exhibiting the collective phenomena of interest: bridging-mediated clustering, competition between clusters, arrested coarsening, and meaningful transcriptional switching statistics. In practice, this proved difficult. Several parameter sets that initially seemed plausible either produced dynamics too weak to generate robust clustering, or instead drove the system into regimes that were too kinetically frozen to be useful for studying variability, commitment, and memory. Establishing the final parameter set therefore required repeated cycles of simulation, visual inspection, quantitative analysis, and revision. In some cases this meant allowing large batches of simulations to run for weeks before concluding that the resulting data were not suitable and that the model had to be retuned and rerun. Although this was one of the most time-consuming parts of the project, it was also one of the most important, because it forced me to engage critically with how microscopic interaction strengths, switching rates, particle numbers, and spatial scales collectively determine the behaviour of the model. The final simulations presented in the dissertation are therefore not based on arbitrary choices, but on a deliberate attempt to work within a specific regime in which fossilisation, commitment, and canalisation could be studied in a controlled way.

By the beginning of the third month, I had started building an analysis pipeline for the simulation output. This initially focused on the static simulations, where the chromatin

state remained fixed throughout the run, and extracted quantities such as burst statistics, correlation times, and related switching observables. Writing this analysis code required its own set of methodological decisions. For instance, I had to decide how TU activity should be defined from the three-dimensional trajectories and whether a telegraph description was an appropriate coarse-grained model for the switching behaviour observed in the simulations. At the same time, it became increasingly clear that a comparison of fixed endpoints was not enough to address the Waddington picture properly, since differentiation is inherently a process rather than just a contrast between pre-defined states.

This led me to the most technically difficult stage of the project: writing new LAMMPS code to simulate the process of "rolling down Waddington's landscape". I had no pre-existing implementation to build directly from, so much of this had to be developed almost from scratch. Here again, the central difficulty was not just coding but deciding what assumptions would best serve the aims of the project. I chose to implement gradual differentiation protocols rather than abrupt state changes, because I wanted to study not only the final differentiated states but also the route by which they were reached, and eventually whether that route could be reversed. This required new LAMMPS routines for recolouring chromatin beads and transcription factors dynamically. Only after this dynamic framework was in place did I begin to think more carefully about quantities such as heterogeneity and plasticity, and about how to interpret them without assigning them an overly literal biological meaning. This part of the project made it clear that analysis was not simply a matter of extracting quantities from the data, but of deciding which observables genuinely captured the physical questions I wanted to ask.

By the fourth month, I had completed the dynamic simulation code and began analysing the resulting trajectories. This involved comparing different differentiation protocols, in particular chromatin remodelling and TF colour collapse, which were loosely inspired by different biological routes to differentiation. I also expanded the Python analysis pipeline to deal with non-stationary simulations through sliding-window techniques. This introduced further trade-offs that had to be considered carefully: the windows needed to be long enough to contain sufficient switching events for the inferred rates to be meaningful, but short enough to follow the gradual temporal evolution of the system. As these results developed, it became increasingly clear that the most interesting aspect of the dynamics was not simply differentiation itself, but the extent to which the differentiated state was remembered. This led naturally to the final stage of the project, which focused on canalisation.

In the final month, I examined how robustly differentiated transcriptional states were retained under reversal protocols and how this depended on variables such as transcription-factor abundance. This allowed me to interpret the valleys of Waddington's landscape as stable attractors whose depth, in the model, increases with TF number. At the same time, this final stage reinforced the importance of interpreting all conclusions within the limits of the specific modelling regime I had constructed. I did not treat the simulations as a literal description of real cell differentiation. Instead, I tried throughout to use them as a controlled physical framework in which assumptions, simplifications, and parameter choices were made deliberately, and in which the validity and consequences of those choices had to be questioned alongside the results themselves.

Acknowledgments

I would first like to acknowledge the School of Physics and Astronomy at the University of Edinburgh for allowing me access to the resources necessary to complete this project. I would also like to thank my supervisor, Prof. Davide Marenduzzo, whose advice has been central to the development of this project. His guidance helped me both to avoid and to overcome obstacles throughout the project. His enthusiasm for the research kept me inspired and engaged throughout. Finally, the countless other friends and colleagues of mine with whom I discussed various aspects of this project have allowed me to structure and order my work in such a way that it is understandable to those not directly involved in the field.

Contents

1 Introduction

Every cell in a multicellular organism contains the same genetic information, yet cells adopt a wide variety of stable identities and specialised roles. Neurons, muscle cells, epithelial cells, and blood cells differ dramatically in structure and function despite sharing an identical genome [1, 2]. Moreover, once a cell has committed to a particular identity, that identity is typically maintained through many rounds of cell division. This stability is essential for the reliable function of tissues and organs. In this sense, cells possess a form of biological memory [3], ensuring that skin cells remain skin cells and do not spontaneously convert into blood or nerve cells.

A useful conceptual framework for understanding this is Waddington's epigenetic landscape [4], in which cellular development is represented as motion across an abstract landscape whose valleys correspond to stable cellular identities. Differentiation, the process through which a cell specialises, may then be understood as a progression from relatively shallow valleys higher in the landscape to more stable, committed states in deeper valleys. In this picture, the central question is what physical and biochemical factors determine the shape of the landscape, and hence the stability of a given cell fate.

Historically, explanations for cell-fate stability have focused primarily on biochemical and molecular mechanisms, especially epigenetic regulation [5]. Chemical modifications of DNA and histone proteins, such as methylation and acetylation, influence which genes are expressed and which are silenced. Here, *gene expression* refers to the process by which the information encoded in a gene is used to produce a functional product, usually a protein. This begins with *transcription*, in which the DNA sequence of a gene is copied into RNA by molecular machinery such as RNA polymerase, and is followed by *translation*, in which that RNA is used as a template to assemble a protein.

What differs from one cell type to another is therefore not the underlying genetic code, but the epigenetic marks imposed upon it. These epigenetic marks help explain how cells maintain long-term memory of identity without altering the DNA sequence itself, since such marks can be inherited across cell divisions [3]. However, epigenetic marks alone do not fully specify how genes are regulated within the nucleus [6, 7].

It has become increasingly clear that cellular identity is also closely linked to the three-dimensional organisation of chromatin within the nucleus [6, 8, 9]. The genome is not a passive linear string, but a folded polymer whose spatial arrangement influences which regulatory elements can interact, and thereby contributes to which genes are stably expressed or silenced. Experimental advances, particularly chromosome conformation capture techniques, have shown that chromatin architecture changes with developmental state. Such changes are especially evident in two stages of cellular development relevant to this project: differentiation [10, 11] and senescence [12, 13]. These observations suggest that chromatin organisation may actively contribute to stabilising or reshaping cellular identity.

Despite this, chromatin folding has often been treated as a technically complex addition to more established epigenetic models rather than as a central organising principle in its own

right. One reason is that the physical mechanism linking chromatin structure to transcriptional regulation has been difficult to formulate within a simple and predictive framework. Recent theoretical work has shown, however, that aspects of chromatin organisation can be understood using ideas from soft matter physics and statistical mechanics. In particular, *bridging-induced phase separation* (BIPS) [14, 15, 16] provides a mechanism by which DNA-binding proteins can drive clustering and large-scale chromatin organisation.

The aim of this project is to use polymer models of BIPS to investigate how changes in chromatin organisation related to stages of cellular development reshape transcriptional dynamics, and how those dynamics can be interpreted through Waddington's landscape. Our focus is not on modelling specific biological cell types, but on understanding how idealised chromatin states associated with different stages of cellular development differ in their transcriptional behaviour. In this sense, the project is concerned less with distinguishing mature identities such as neurons or epithelial cells, and more with understanding the physical changes that accompany movement between developmental stages.

Rather than attempting to model the full biochemical complexity of real chromatin, we work with simplified stochastic "toy" models that retain the ingredients most relevant to the present question: chromatin as a polymer [17, 18, 19], transcription factors as diffusing DNA-binding proteins [20, 21], and transcriptional activity as a stochastic ON/OFF process [22]. By constructing chromatin models corresponding to different stages of cellular development, we ask how changes in three-dimensional organisation alter patterns of transcription.

To quantify these transcriptional dynamics, we analyse transcription-unit activity using the telegraph model [23, 24], extracting effective switching rates and burst statistics from simulated time series. These quantities provide a compact and physically interpretable summary of transcriptional behaviour. By comparing them across different forms of chromatin organisation, we show that the structural and dynamical consequences of BIPS can be related to concepts such as fossilisation, developmental commitment, and cellular memory.

2 Background

2.1 Waddington's Landscape, Cellular Development, and Chromatin Organisation

Waddington's epigenetic landscape provides a useful conceptual framework for understanding cellular identity [4]. In this picture, the current state of a cell as it progresses through development is represented by a marble moving across an abstract landscape, while stable cell fates, corresponding to the final specialised identities of the cell, are represented by valleys in that landscape, as illustrated in Figure 1. Differentiation may then be viewed as a progression into deeper valleys, with increasingly specialised cell states becoming more stable and resistant to perturbation. The central idea relevant to the present work is therefore not the literal geometry of the landscape, but the notion that some cellular states are more plastic, while others are more robust and effectively locked in.

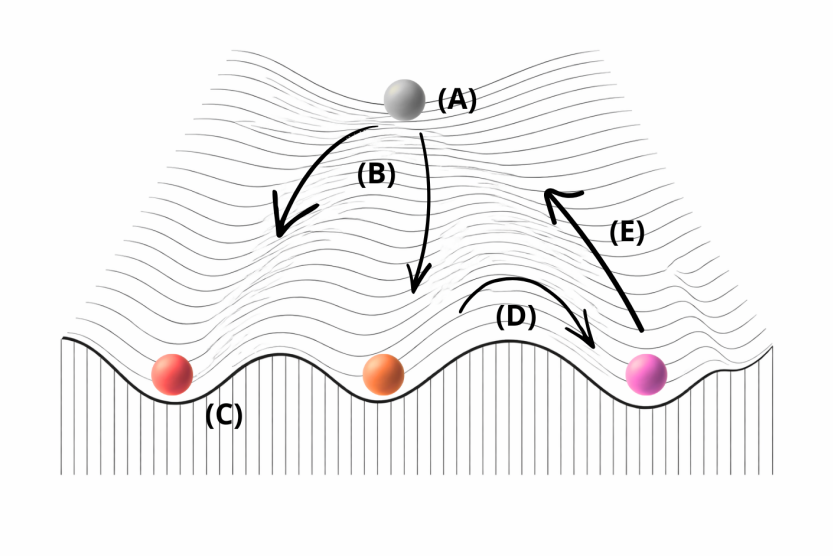

Figure 1. Schematic illustration of Waddington's epigenetic landscape. Spheres represent the instantaneous regulatory state of a cell. **(A)**: The cell begins as a stem cell, which has no specialised role and can adopt many identities. **(B)**: Differentiation is associated with transitions into deeper valleys. **(C)**: These deeper valleys represent stable cellular identities. **(D)**: The possibility of changes in cellular identity due to fluctuations (transdifferentiation). **(E)**: The ability of cells to be returned to a stem-cell state. In this project, chromatin organisation is proposed to influence the shape of the landscape.

This picture has become particularly powerful when combined with the stochastic nature of gene regulation. Because gene expression and molecular interactions are noisy, cells do not sit at a single fixed point in the landscape. Rather, they fluctuate around relatively stable states, and under sufficiently strong perturbations or over long timescales can in principle transition between them [25]. The landscape therefore provides a language for discussing both stability, through valley depth, and variability, through fluctuations within a valley. This makes it especially useful for the present project, where the simulated chromatin–TF dynamics are intrinsically stochastic.

A central question is what determines the shape of this effective landscape in real cells. Biochemical and epigenetic factors clearly play a major role, but the genome is not simply a linear sequence of information: in vivo it is a three-dimensional object whose organisation is closely linked to gene regulation [6]. Developments in chromosome conformation capture and related approaches have established that chromatin architecture changes with developmental state, and that three-dimensional organisation can influence which genes are expressed [8]. From the perspective of Waddington's landscape, chromatin organisation can therefore be interpreted as one of the factors helping to shape the depth, accessibility, and stability of different valleys.

This framework is useful for the three stages of cellular development studied in this project: a

stem-like state, a differentiated-like state, and a senescent-like state. The word "like" is used deliberately, since our models are idealised toy models rather than direct representations of specific biological cell types. A *stem cell* retains both the ability to self-renew and the ability to differentiate into multiple specialised lineages. In the language of the landscape, it therefore occupies a relatively shallow region higher up in the landscape in which multiple future fates remain accessible. In terms of chromatin organisation, this motivates a state with relatively high configurational freedom and a broad range of possible transcriptional outcomes.

A *differentiated cell*, by contrast, has adopted a specialised function and typically displays more constrained gene expression. Such cells correspond to deeper valleys in the landscape, reflecting greater stability of cellular identity and reduced accessibility of alternative fates. Within the present project, this motivates chromatin states in which transcriptional possibilities are more restricted and a smaller subset of regulatory patterns remains dynamically available.

Finally, a *senescent cell* is a cell that has exited the cell cycle irreversibly in response to stress or damage and undergone substantial architectural reorganisation. Senescence is generally regarded as a terminal state with altered chromatin structure and reduced capacity for further chromatin reconfiguration [26, 27, 28]. In the language of the landscape, it is therefore natural to regard senescence as a strongly stabilised and structurally constrained state. We include it here because its associated chromatin organisation is comparatively well characterised and provides a useful contrast to differentiation.

There is substantial experimental evidence that chromatin organisation varies across these stages of cellular development, although the precise nature of the reorganisation depends on both the spatial scale considered and the biological context. During stem-cell differentiation, large-scale architectural features such as topologically associating domains (TADs) can remain relatively stable, while substantial reorganisation occurs within and between domains, including widespread switching of A/B compartments and changes in local interaction structure [11]. By contrast, senescent cells often undergo more pronounced large-scale chromatin reorganisation, frequently including the formation of senescence-associated heterochromatin foci (SAHF) in oncogene-induced senescence, although the details depend on the form of senescence, the inducing stress, and the cell type [12, 13]. Taken together, these observations indicate that chromatin organisation is not fixed, but varies significantly with developmental state.

Rather than assuming *a priori* how such changes map onto deeper or shallower valleys in the landscape, a central aim of this work is to model these stages by deliberately varying chromatin architecture and then asking how the resulting transcriptional behaviour changes. The underlying hypothesis is that if chromatin organisation helps shape the effective landscape, then altering chromatin structure should change the transcriptional trajectories accessible to the system. In other words, changing how chromatin can fold and cluster should change which parts of the genome are preferentially engaged by transcription factors, and hence should reshape the resulting transcriptional programme.

This motivates the next step: introducing a physical mechanism that can relate chromatin organisation to transcriptional activity in a minimal and explicit way. In the following subsection, we therefore turn to bridging-induced phase separation, which provides a concrete route by which diffusing DNA-binding proteins can reorganise chromatin and generate transcriptionally active clusters.

2.2 Bridging-Induced Phase Separation

If chromatin organisation helps shape the effective developmental landscape, the next question is what physical mechanism can relate chromatin structure to transcriptional activity. To address this, we first consider the broader idea of phase separation. Over the past decade, phase separation has emerged as a possible organising principle for many biological processes, including the formation of dynamic intracellular compartments [29]. In the nucleus, where chromatin — the DNA–protein complex that packages the genome — is highly abundant, this has led to the suggestion that aspects of chromatin organisation may also be understood in terms of phase-separation-like phenomena [30, 31, 32].

A difficulty, however, is that the term "phase separation" is often used rather broadly in biology [29]. In the simplest picture of liquid–liquid phase separation (LLPS), clustering arises because direct attractive interactions between molecules favour demixing into coexisting phases. While this may describe some nuclear compartments, it is not clear that all chromatin-associated clustering can be explained in this way. In particular, chromatin is not simply a collection of freely diffusing molecules: it is a long, connected polymer whose organisation is constrained by chain connectivity, bending stiffness, and steric repulsion between its segments. Any mechanism for chromatin organisation must therefore account for this polymeric structure.

This motivates an alternative mechanism known as *bridging-induced phase separation* (BIPS) [14, 15, 16]. In BIPS, clustering does not require intrinsic attractive interactions between proteins. Instead, it arises because multivalent DNA-binding proteins can simultaneously bind more than one region of the chromatin fibre. When such a protein bridges two distant chromatin segments, it locally increases chromatin density. This in turn makes further binding events in the same region more likely, creating a positive feedback loop between local chromatin density and local protein concentration. Repeated bridging can therefore drive the emergence of locally dense chromatin–protein clusters.

In contrast to LLPS, the key ingredient in BIPS is not direct protein–protein attraction but coupling between protein binding and polymer organisation. Closely related descriptions have framed this as an effective polymer–polymer phase separation, emphasising that it is the chromatin fibre itself that segregates into regions of differing density [33]. For the present project, this is particularly attractive because it provides a direct route by which changes in chromatin organisation could alter transcriptional behaviour without requiring an explicit biochemical model of protein condensates.

The relevance of this idea is supported by experimental observations that bound proteins can assemble into hot spots in the fly genome [34], and that active RNA polymerase II molecules

form transcription-associated clusters or "factories" [35, 36]. These findings sit within a broader body of work in which DNA folding is modelled as being driven by DNA-binding proteins [37, 38, 39, 40, 41, 42]. Many other simulations and experiments are also consistent with clustering and compaction driven by bridging [43, 44, 45, 46, 47].

To make the BIPS mechanism more precise, it is useful to introduce a continuum description. We consider two spatially varying fields: a chromatin density field $\rho(\mathbf{x}, t)$, along which lie specific binding sites that we later model as transcription units (TUs), and a second field $\Phi(\mathbf{x}, t)$ representing the local density of binding transcription factors and RNA polymerases, which we collectively refer to as TFs. The exact microscopic implementation of these interactions is introduced later in the *Simulation and Analysis Methods* section (Sec. 3.1).

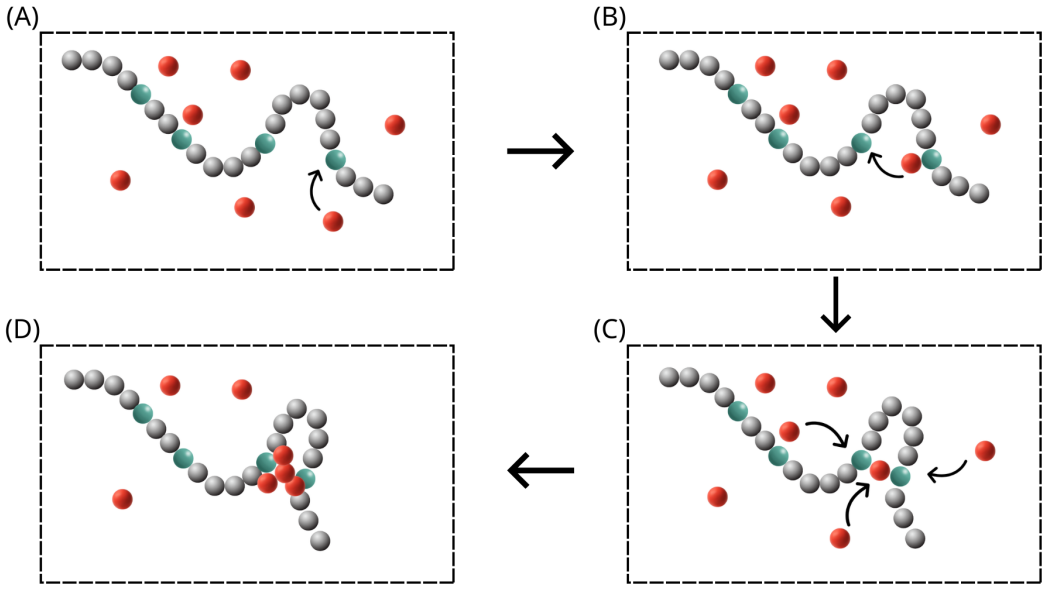

Figure 2. Illustration of bridging-induced phase separation (BIPS) in chromatin. Green beads represent transcription units (TUs), which together with grey beads represent chromatin, while orange beads represent multivalent transcription factors (TFs). **(A)** TFs diffuse and bind transiently to individual TUs along the chromatin fibre. **(B)** A multivalent TF simultaneously binds two distant TUs, creating a molecular bridge that locally increases chromatin density. **(C)** The increase in local chromatin density enhances the probability of further TF binding in the same region, generating a positive feedback loop. **(D)** Repeated bridging events lead to the formation of stable, locally dense clusters of TFs and chromatin. Importantly, TFs do not possess any direct attractive interactions with one another; clustering arises solely from chromatin–TF binding (red arrows) and chromatin flexibility.

The essential BIPS feedback loop can then be expressed as a coupling between these two fields: local enrichment of TFs promotes local chromatin enrichment, because multivalent binding draws distant polymer segments together, while local chromatin enrichment increases the likelihood of further binding in the same region, as illustrated in Figure 2.

A standard way to model phase separation in conserved densities is through Cahn–Hilliard dynamics, with a reaction term for $\Phi(\mathbf{x}, t)$. A minimal coupled form consistent with BIPS can be written as

$$\partial_t \rho = M_\rho \nabla^2 \mu_\rho = M_\rho \nabla^2 \big(a_1 \rho - k \nabla^2 \rho - \chi \Phi + g \rho^3 \big) , \tag{1}$$

$$\partial_t \Phi = M_\Phi \nabla^2 \mu_\Phi - \alpha(\Phi - \Phi_0) = M_\Phi \nabla^2 (a_2 \Phi - \chi \rho) - \alpha(\Phi - \Phi_0), \tag{2}$$

where M_ρ and M_Φ set diffusive mobilities, k penalises sharp spatial gradients in ρ by giving interfaces a finite width, and g stabilises the density at large ρ. The coupling strength χ encodes the feedback between chromatin density and TF concentration. The additional term $-\alpha(\Phi - \Phi_0)$ represents stochastic switching of TFs towards a baseline level Φ_0, reflecting the fact that TFs can switch between binding and non-binding states [15].

Equations (1)–(2) can be derived from an effective free-energy functional $F[\rho, \Phi]$, with chemical potentials $\mu_\rho = \delta F / \delta \rho$ and $\mu_\Phi = \delta F / \delta \Phi$. If we neglect TF switching for the moment, that is, set $\alpha = 0$, a functional consistent with the equations above is

$$F[\rho, \Phi] = \int d^3x \Big[\frac{a_1}{2} \rho^2 + \frac{g}{4} \rho^4 + \frac{k}{2} |\nabla \rho|^2 + \frac{a_2}{2} \Phi^2 - \chi \rho \Phi \Big], \tag{3}$$

which yields $\mu_\rho = a_1 \rho + g \rho^3 - k \nabla^2 \rho - \chi \Phi$ and $\mu_\Phi = a_2 \Phi - \chi \rho$. The coupling term $-\chi \rho \Phi$ makes the positive feedback explicit: increasing Φ lowers the local free-energy cost of increasing ρ, while increasing ρ lowers the cost of increasing Φ [14, 15, 16].

This continuum picture also clarifies what distinguishes BIPS from LLPS. In LLPS, one typically assumes direct intermolecular attractions that drive demixing. In BIPS, clustering can arise even without intrinsic TF–TF attraction, because binding to a polymer generates an effective attraction mediated by the chromatin fibre itself [14]. Multivalency is therefore essential: without the ability to bridge multiple sites, the coupling between ρ and Φ is too weak to sustain clustering [14]. The plausibility of this assumption is supported both by the broader literature [37] and by evidence that many transcription factors are bivalent or multivalent [48].

The switching term in Eq. (2) provides a further important ingredient. If TFs remained permanently in their binding state, phase separation would generally continue to coarsen, producing clusters that grow without bound over time [15]. By allowing TFs to switch stochastically between binding and non-binding states at a characteristic rate α, coarsening can instead be arrested, yielding persistent yet dynamic clusters [15].

This immediately suggests a link to transcription. Since the TFs in our coarse-grained description include transcription factors and RNA polymerase II, BIPS-driven clustering naturally generates regions with enhanced local transcriptional activity — often described in the literature as transcription factories, hubs, condensates, or super-enhancer-associated clusters [49, 50, 51, 52, 53]. In such regions, the local concentration of TFs is elevated, making nearby transcription units more likely to be engaged and hence more likely to be transcriptionally active.

This is precisely the feature that makes BIPS useful for the present work. If different forms of chromatin organisation alter where bridges can form and which clusters emerge, then they should also alter the transcriptional output associated with those clusters. For example, restricting TF access to selected chromatin regions should reduce their participation in bridging and hence modify the resulting transcriptional activity. This provides a concrete mechanism by which different stages of cellular development, interpreted through Waddington's landscape, can be associated with different transcriptional programmes.

The next step is therefore to move from this continuum description to an explicit polymer model in which chromatin is represented as a connected chain and TFs as diffusing multivalent binders. This will allow us to study, in a controlled way, how changes in chromatin organisation reshape clustering, transcriptional dynamics, and ultimately the effective stability of different cell states.

3 Simulation and Analysis Methods

3.1 Simulating Chromatin and Transcription Factors

In this section, we describe the coarse-grained polymer model used to simulate chromatin organisation and transcription factor (TF) clustering in three dimensions. The model is designed to realise bridging-induced phase separation (BIPS), in which multivalent binders bridge distant chromatin sites and thereby generate a positive feedback between local chromatin density and local TF concentration. Our implementation builds on the pangenomic framework introduced in Ref. [54] and related work [15, 16], but extends it by introducing *five* distinct classes of transcription units (TUs), each with its own cognate TF partner. This multicolour extension replaces the simpler regulatory scheme of earlier models with a richer regulatory landscape, allowing differentiation to be modelled as a reduction in regulatory degrees of freedom. In Sec. 3.2, we further introduce explicit dynamical protocols for differentiation and senescence in order to study how chromatin remodelling and other forms of regulatory change reshape transcriptional dynamics. The model retains only the ingredients required for selective binding, bridging, clustering, and stochastic switching, providing a minimal framework in which commitment and memory can be studied in a physically interpretable way.

Chromatin Beads and Transcription Factors

We model the chromatin fibre as a bead–spring polymer in a three-dimensional periodic box of size $L_x \times L_y \times L_z$, with periodic boundary conditions in all directions. Each polymer bead represents a coarse-grained chromatin segment of approximately 3 kbp (kilo–base-pairs), while TFs are represented as diffusing spherical particles that can exist in either a *binding* or *non-binding* state.

In the coarse-grained model, chromatin beads and TFs are assigned the same effective bead diameter in reduced Lennard–Jones units. This does not imply that they are biologically comparable in size. A chromatin bead represents a DNA segment of order 3 kbp, whereas an individual TF typically binds a DNA region of only 10–20 bp [55], so the underlying biological scales differ by roughly two orders of magnitude. The common bead diameter should therefore be understood as an effective interaction length scale, introduced to provide a minimal description of steric effects, TF–chromatin binding, and clustering, rather than as a literal molecular radius. Likewise, the model does not attempt to resolve the biochemical details of transcription directly, but instead retains only the ingredients needed to study selective binding, bridging, clustering, and stochastic switching.

Binding specificity is encoded by assigning discrete particle types to both TFs and chromatin beads. TFs are represented by ten particle types in total, organised into five species. For each species, which we informally refer to as a "colour" $c \in \{1, \ldots, 5\}$, the odd-labelled type $(2c-1)$ denotes the *binding* state, while the even-labelled type $(2c)$ denotes the corresponding *non-binding* state, so that each TF colour is represented by a binding and a non-binding version. This explicit separation allows TF switching to be implemented directly by changing particle type.

Chromatin beads are assigned seven types in total. Type-11 beads are inert and do not participate in attractive interactions. The most numerous chromatin beads are type 12, which act as non-specific binding sites. Finally, beads of types 13–17 represent five distinct classes of transcription units (TUs). Each TU colour $c \in \{1, \ldots, 5\}$ is associated with chromatin bead type $(12 + c)$ and is designed to interact selectively with the binding TF of the same colour. The full type assignment used throughout the simulations is summarised in Table 1.

Table 1. Summary of bead types used in the coarse-grained chromatin–TF model. Here $c \in \{1, \ldots, 5\}$ labels the TF/TU colour.

Type(s)	**Class**	**Identity**	**Role**
$2c - 1$	TF	colour c, binding	binds type 12 weakly and type $12 + c$ strongly
$2c$	TF	colour c, non-binding	non-binding version of TF colour c
11	chromatin	inert bead	neutral background; no attractive interactions
12	chromatin	non-specific site	weakly binds all binding TF species
$12 + c$	chromatin	TU colour c	strongly binds the binding TF of colour c

This type-based description defines the degrees of freedom of the model without yet specifying how the interactions are implemented. In the following subsection, we show how these particle identities are translated into steric, non-specific, and specific binding interactions through a common Lennard–Jones framework.

Interaction rules and binding specificity

Having defined the particle types, we now specify how they interact. All non-bonded interactions in the model are implemented through the same truncated Lennard–Jones (LJ) pair potential acting between particles i and j separated by a distance r_{ij},

$$U_{\mathrm{LJ}}(r_{ij}) = 4\varepsilon_{ij}\left[\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\right], \qquad U_{\mathrm{pair}}(r_{ij}) = \begin{cases} U_{\mathrm{LJ}}(r_{ij}), & r_{ij} < r_{ij}^{(c)}, \\ 0, & r_{ij} \geq r_{ij}^{(c)}. \end{cases} \tag{4}$$

Here σ_{ij} sets the effective interaction diameter, ε_{ij} the interaction strength, and $r_{ij}^{(c)}$ the cut-off radius. Throughout this work we use reduced Lennard–Jones units and set $\sigma_{ij} = 1$ for all interacting pairs. Different classes of interaction are then specified by choosing appropriate values of ε_{ij} and $r_{ij}^{(c)}$.

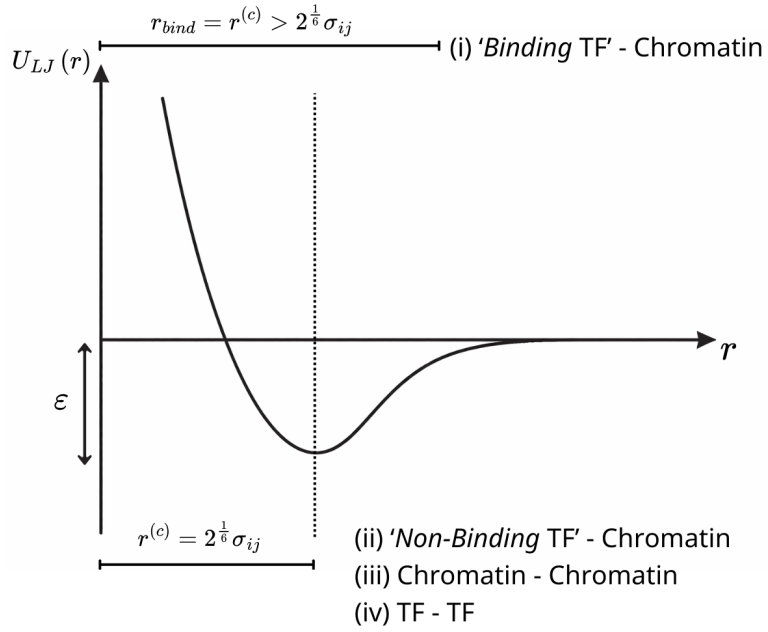

Figure 3. Schematic Lennard–Jones potential illustrating how different interaction classes are implemented in the model. For attractive binding TF–chromatin pairs, the cutoff is extended beyond the potential minimum, $r_{\mathrm{bind}} = r^{(c)} > 2^{1/6}\sigma_{ij}$, so that both the repulsive core and attractive well are retained. By contrast, for steric-only pairs, including non-binding TF–chromatin, chromatin–chromatin and TF–TF interactions, the cutoff is set at the Lennard–Jones minimum, $r^{(c)} = 2^{1/6}\sigma_{ij}$, so that only excluded-volume repulsion remains.

For steric-only pairs, the cut-off is placed at the Lennard–Jones minimum,

$$r_{ij}^{(c)} = 2^{1/6}\sigma_{ij}, \tag{5}$$

so that the attractive tail is removed, leaving a short-ranged excluded-volume interaction. In the present model, this purely repulsive form applies to chromatin–chromatin pairs, non-binding TF–chromatin pairs, and all TF–TF pairs. Steric exclusion between all particles

prevents unphysical overlap, while ensuring that any attractive association arises only from the selected binding interactions introduced below.

Attractive TF–chromatin binding is encoded by extending the cut-off beyond the minimum of the LJ potential, so that both the repulsive core and attractive well are retained. Binding is introduced only for binding TFs; non-binding TFs do not participate in any attractive TF–chromatin interactions. In this way, TF binding is not imposed through a separate potential, but is built directly into the same pairwise LJ framework used for steric exclusion. A summary of how the truncated Lennard–Jones potential is used for attractive and steric-only interaction classes is shown in Figure 3.

Two distinct binding modes are included in the model: weak non-specific binding and strong specific binding. Non-specific binding is assigned to type-12 chromatin beads, which act as low-affinity binding sites distributed along the polymer. All binding TF species interact attractively with these beads according to

$$\varepsilon_{(2c-1),\,12} = \varepsilon_{\mathrm{ns}}, \qquad r^{(c)}_{(2c-1),\,12} = r_{\mathrm{bind}} > 2^{1/6}\sigma_{ij}, \qquad c \in \{1,\ldots,5\}, \tag{6}$$

where $\varepsilon_{\mathrm{ns}}$ is the common non-specific interaction strength and r_{bind} is the common binding cut-off. This interaction provides a baseline level of chromatin association for all binding TF colours, allowing them to sample the fibre and participate in bridging even away from transcription units.

Specific binding is implemented through the five TU bead types 13–17. Each TU colour recruits only the binding TF of the same colour (called its "cognate" TF) through a stronger attractive interaction,

$$\varepsilon_{(2c-1),\,(12+c)} = \varepsilon_{\mathrm{sp}}, \qquad r^{(c)}_{(2c-1),\,(12+c)} = r_{\mathrm{bind}} > 2^{1/6}\sigma_{ij}, \qquad c \in \{1,\ldots,5\}, \tag{7}$$

with $\varepsilon_{\mathrm{sp}} > \varepsilon_{\mathrm{ns}}$. The allowed specific interactions are therefore

$$(1 \leftrightarrow 13),\ \ (3 \leftrightarrow 14),\ \ (5 \leftrightarrow 15),\ \ (7 \leftrightarrow 16),\ \ (9 \leftrightarrow 17),$$

where the first entry denotes the binding TF type and the second the corresponding TU bead type. Non-cognate TF–TU pairs are not assigned this stronger attraction. Thus, most chromatin beads provide weak non-specific binding, while a sparse subset of TU beads act as strong, selective recruitment sites for their cognate TF species. These rules are illustrated schematically in Figure 4.

Throughout this work, we use a common short-ranged binding radius r_{bind} for both non-specific and specific attractive interactions. This keeps binding local in three-dimensional space while still allowing multivalent bridges to form often enough for clustering and switching statistics to emerge clearly. Choosing $\varepsilon_{\mathrm{ns}} < \varepsilon_{\mathrm{sp}}$ establishes a hierarchy of affinities in which TU beads act as strong nucleation sites for recruitment and bridging, while non-specific binding remains sufficient to maintain generic TF–chromatin association along the polymer.

These interaction rules are what allow binding TFs to form bridges between distinct chromatin segments. Because each binding TF can interact attractively with more than one

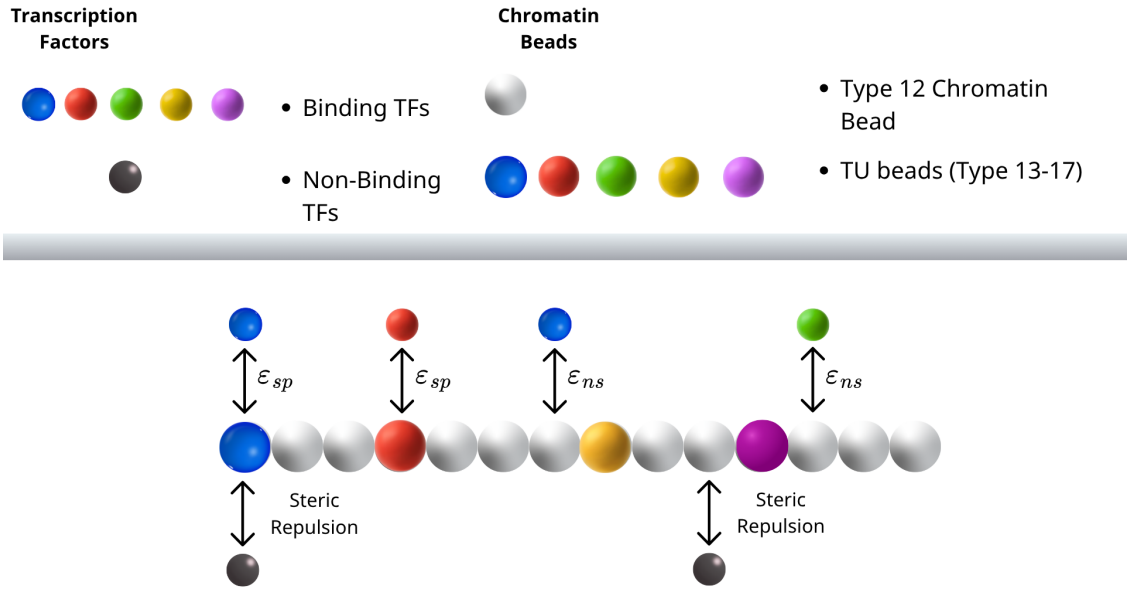

Figure 4. Schematic illustration of the binding rules used in the multicolour chromatin–TF model. Binding TFs are shown as coloured particles, while non-binding TFs are shown in black. Along the chromatin fibre, most beads are type-12 beads (grey), which act as weak non-specific binding sites for all binding TF species with interaction strength $\varepsilon_{\mathrm{ns}}$. A sparse subset of beads are transcription units (TU beads; types 13–17), shown in colour, which bind strongly and selectively to their cognate binding TF species with interaction strength $\varepsilon_{\mathrm{sp}}$. Non-binding TFs do not participate in attractive interactions and interact with chromatin only through steric repulsion.

chromatin bead, local accumulation of TFs promotes additional local binding and bridging, generating the positive feedback characteristic of BIPS. Crucially, the model contains no direct TF–TF attraction: any clustering arises solely through chromatin-mediated bridging. This deliberately isolates the BIPS mechanism and makes the physical origin of the resulting clusters easier to interpret.

Polymer connectivity and stochastic dynamics

Having specified the non-bonded interaction rules, we now describe how the chromatin fibre is bonded together and how the full system is evolved in time. The chromatin fibre is modelled as a semiflexible self-avoiding bead–spring polymer, whose structure is determined by chain connectivity, bending rigidity, and the non-bonded steric and binding interactions described above. These ingredients are illustrated schematically in Figure 5.

Polymer connectivity is enforced by bonding neighbouring chromatin beads along the chain. In our simulations, these bonds are modelled using finitely extensible non-linear elastic (FENE) springs,

$$U^{\mathrm{FENE}}(r_{ij}) = -\frac{K R_0^2}{2} \ln\left[1 - \left(\frac{r_{ij}}{R_0}\right)^2\right], \tag{8}$$

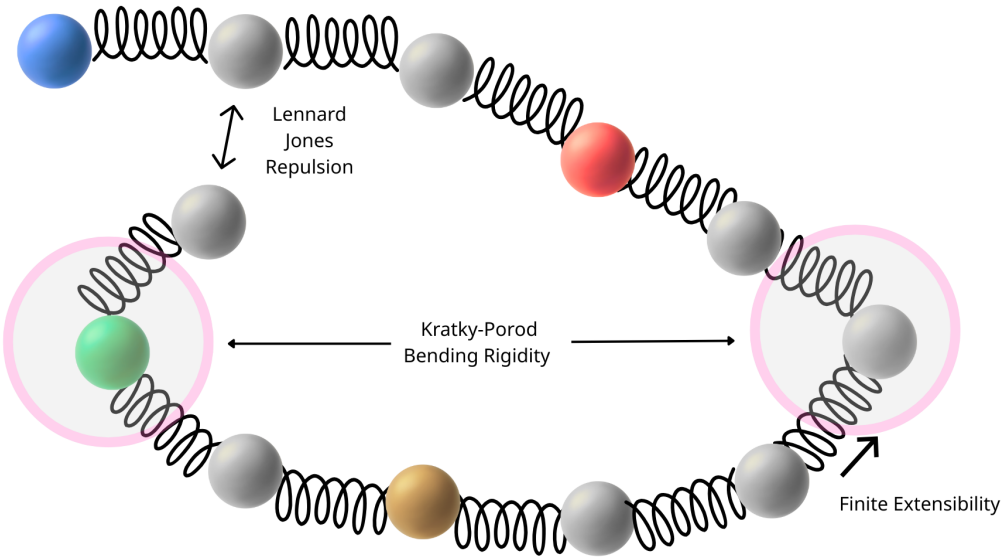

Figure 5. Schematic of the bead–spring polymer mechanics used to model chromatin. Neighbouring beads are connected by finitely extensible non-linear elastic (FENE) bonds, which preserve chain connectivity while preventing unphysical bond stretching. Excluded-volume interactions are enforced via a Lennard–Jones potential truncated at its minimum, producing short-ranged steric repulsion. A Kratky–Porod angle potential introduces bending rigidity by penalising sharp local bends, thereby giving the polymer finite stiffness.

where K sets the bond stiffness and R_0 is the maximum allowed bond extension. This potential permits thermal fluctuations in bond length while strongly penalising overstretching, thereby maintaining polymer integrity throughout the simulation.

To capture the semiflexible nature of chromatin, we also include bending rigidity through a three-body angle potential acting on successive bead triplets $(i, i+1, i+2)$,

$$U_{\text{bend}}(\theta_i) = K_\theta \left(1 + \cos\theta_i\right), \tag{9}$$

where θ_i is the angle between adjacent bond vectors. This is the standard Kratky–Porod form for a worm-like chain, in which the bending stiffness K_θ controls the persistence length of the polymer: increasing K_θ suppresses sharp bends and produces a stiffer fibre.

The time evolution of all particles, both chromatin beads and TFs, is simulated using Langevin dynamics. For each particle i with position $\mathbf{r}_i(t)$ and mass m_i,

$$m_i \ddot{\mathbf{r}}_i = -\nabla_i U - \gamma \dot{\mathbf{r}}_i + \sqrt{2\gamma k_{\text{B}} T}\, \boldsymbol{\eta}_i(t), \tag{10}$$

where U is the total potential energy, including bond, bending, and pairwise interaction terms, γ is a friction coefficient, and $\boldsymbol{\eta}_i(t)$ is a Gaussian white-noise term with zero mean and unit variance. In practice, the dynamics are overdamped on the timescales of interest. We work throughout in reduced Lennard–Jones units, with $\sigma = \varepsilon = m = 1$, at fixed

temperature T. This is important because all interaction parameters in the model are defined relative to the thermal scale $k_{\mathrm{B}}T$: binding affinities, steric repulsions, and bending rigidities influence the dynamics only through their competition with thermal fluctuations. The Langevin thermostat therefore acts as a heat bath, balancing dissipative losses with stochastic forcing so that the system remains at the chosen target temperature throughout the simulation.

To avoid numerical artefacts arising from initial bead overlaps and sharp kinks in the randomly generated starting configurations, we employ a short initial equilibration stage before the production dynamics begin. During this relaxation phase, steric interactions are implemented using a soft repulsive potential, polymer bonds are represented by harmonic springs, and a relatively large bending stiffness is applied. This allows the initial chain to relax smoothly, removing overlaps and strongly bent local configurations without generating excessively large forces or artificial bond stretching. After this preliminary equilibration, the target simulation conditions are restored: harmonic bonds are replaced by FENE bonds, the bending stiffness is reduced to its intended value, and the full Lennard–Jones interaction matrix implementing steric, non-specific, and specific TF–chromatin interactions is activated.

TF switching

The ingredients introduced above are sufficient to generate chromatin bridging and local TF clustering, but an additional non-equilibrium process is required to prevent these clusters from coarsening indefinitely. Following earlier work on ephemeral protein binding in BIPS models [15], we therefore allow TFs to switch stochastically between *binding* and *non-binding* states. This continual turnover between binding and non-binding states is a key ingredient in producing arrested coarsening, so that clusters remain dynamic rather than collapsing irreversibly into a small number of large, frozen aggregates.

In the simulation, switching is implemented by changing the particle type of a TF between the two states associated with its colour. For each colour $c \in \{1, \ldots, 5\}$, the binding TF is type $(2c-1)$ and the corresponding non-binding TF is type $(2c)$. A switching event therefore consists of reassigning particles of type $(2c-1)$ to type $(2c)$, and vice versa, while preserving the total number of TFs of each colour.

Switching is applied at discrete intervals Δt_{sw} during the simulation. At each switching event, a fraction f of TF particles in each colour is randomly reassigned from binding to non-binding and from non-binding to binding. This discrete update rule approximates a Poisson switching process with an effective rate

$$\alpha \sim \frac{f}{\Delta t_{\mathrm{sw}}}, \tag{11}$$

in simulation time units, up to discretisation effects.

In practice, switching is applied every 10^4 timesteps, with a fraction $f = 0.01$ of each TF colour reassigned between its binding and non-binding states. These parameter values are

chosen to match the switching regime used in earlier BIPS studies [15]. They ensure that TF switching acts on a timescale slower than local chromatin–TF rearrangements, so that binding and unbinding events can still generate bridging and local structure, while remaining fast enough to prevent irreversible growth of a few dominant clusters.

Initial conditions and ensemble of independent runs

Initial configurations are generated using a dedicated initialisation routine. TFs are placed uniformly at random in the periodic box, corresponding to a well-mixed starting state, while the chromatin fibre is generated as a confined three-dimensional random walk within the same box and then relaxed under the equilibration protocol described above.

In all simulations, the system contains a single polymer of $N = 6000$ beads in a periodic box of size $100 \times 100 \times 100$, together with TF particles of types 1–10. Since these ten types represent five TF species, each split into binding and non-binding states, the baseline choice $n_1 = \cdots = n_{10} = 40$ corresponds to 40 particles per bead type, or 80 particles per TF species when the two states are combined. These choices place the system in a sufficiently dense yet spatially extended regime to permit non-trivial long-range folding, multiple interacting chromatin clusters, and competition between those clusters.

Transcription units are assigned stochastically along the polymer: each bead is designated a TU with probability p_{TU}, corresponding to a mean spacing $\ell_{\mathrm{TU}} \sim 1/p_{\mathrm{TU}}$ beads between TUs. Each TU bead is then assigned one of five colours with equal probability, thereby defining its cognate TF partner through Eq. (7). The mean TU spacing was chosen to avoid the trivial limits of an almost uniformly active fibre on the one hand and an overly sparse regulatory landscape on the other, thereby placing the system in a regime where competition between distinct active sites can emerge.

Table 2 summarises the core interaction and dynamical parameters in reduced LJ units. These values were chosen to place the system in a regime where clusters form but remain dynamic, so that binding competition, cluster rearrangement, and arrested coarsening can all be captured on accessible simulation timescales. In exploratory runs, weaker interactions produced little sustained clustering, whereas substantially stronger interactions led to overly stable aggregates and strongly suppressed switching. The parameter set used here was therefore selected as a compromise that preserves both robust bridge-induced clustering and ongoing transcriptional dynamics. The aim was not to perform an exhaustive scan of parameter space, but rather to work in a representative regime in which the consequences of changing chromatin organisation could be meaningfully resolved.

Because the model is intrinsically stochastic, all observables are obtained from an ensemble of independent simulations started from distinct initial conditions and random seeds. In practice, we analyse 50 such runs and report statistics averaged across the ensemble.

Table 2. Core simulation parameters in reduced LJ units. Values shown correspond to the baseline "stem-like" simulations; modified values for differentiation and senescence protocols are described in Sec. 3.2.

Quantity	Value / description
Units	Lennard–Jones: $\sigma = \varepsilon = m = 1$
Boundary conditions	Periodic, $L_x = L_y = L_z = 100$
Polymer bond potential	FENE, Eq. (8), $K = 30.0$, $R_0 = 1.6$
Polymer bending	Cosine, Eq. (9), $K_\theta = 3.0$
Sterics	Repulsive LJ, Eq. (4), with $r_c = 2^{1/6}\sigma \approx 1.122$
Non-specific binding	$\varepsilon_{\mathrm{ns}} = 3.0$ for binding TF–type 12, Eq. (6)
Specific binding	$\varepsilon_{\mathrm{sp}} = 8.0$ for cognate pairs, Eq. (7)
Binding range	$r_{\mathrm{bind}} = 1.8$ for attractive pairs
Dynamics	Langevin, Eq. (10), $T = 1.0$, $\gamma = 1.0$
Time step	$\delta t = 0.01$
Switching interval	$\Delta t_{\mathrm{sw}} = 10^4$ timesteps
Switching fraction	$f = 0.01$ per event
Independent runs	50
Mean TU spacing	$\ell_{\mathrm{TU}} = 30$ beads

Temporal scale of the simulations

It is useful at this stage to clarify the temporal scales resolved by the simulations, since this sets the context for everything that follows. In particular, when differentiation and senescence are later introduced as dynamical processes, and memory is studied through reversal and hysteresis, it is important to understand the timescales on which these perturbations are imposed relative to the underlying chromatin–TF dynamics.

Throughout this work, the equations of motion are integrated using an elementary timestep of $\delta t = 0.01$ in simulation time units derived from Lennard–Jones units. The saved trajectory frames used for later analysis are, however, much more widely separated than this microscopic update scale. In the baseline simulations, frames are written every 10^5 integration timesteps, corresponding to 10^3 simulation time units between successive snapshots. A saved frame should therefore not be interpreted as an elementary update of the system, but as a coarse-grained snapshot taken after substantial chromatin and TF motion has already occurred. This coarse sampling reduces strong frame-to-frame correlations and keeps the output manageable, while still resolving the slower evolution of clustering and transcriptional activity relevant to the present work.

This temporal structure becomes clearer when compared with the TF switching dynamics. As described above, switching events are applied every 10^4 integration timesteps, corresponding to approximately 10^2 simulation time units. Because only a fraction $f = 0.01$ of each TF species is reassigned at any given event, the effective switching timescale of an individual TF remains longer than the update interval itself. The simulations therefore resolve dynamics on three distinct temporal levels, summarised in Table 3: rapid local rearrangements of

Table 3. Summary of the three main temporal scales resolved in the simulations.

Level	**Timescale**	**Physical significance**
Integration timestep	$\delta t = 0.01$ simulation time units	Elementary Langevin update scale, resolving rapid local motion of chromatin beads and TF particles.
TF switching update	10^4 integration timesteps = 10^2 simulation time units	Intermediate timescale on which TF identities are reassigned, controlling stochastic binding/non-binding switching and recolouring.
Saved trajectory frame	10^5 integration timesteps = 10^3 simulation time units	Coarse-grained sampling scale used in the analysis, over which cumulative changes in cluster organisation and TU activity are observed.

chromatin–TF contacts on the scale of individual integration steps, repeated population-level TF switching updates on intermediate timescales, and slower cumulative changes in cluster structure and transcriptional activity over the interval between saved frames.

This separation of timescales is important for interpreting the later results. Even trajectories containing only around 10^2 saved frames already sample a substantial amount of underlying dynamics, including many switching updates and repeated opportunities for cluster rearrangement. It also means that when differentiation is imposed gradually through the dynamical protocols introduced below, it occurs on a timescale that can be compared directly with the switching and rearrangement times already present in the system, rather than as an effectively instantaneous perturbation. Understanding this temporal structure is therefore essential both for interpreting the transcriptional activity traces extracted from the simulations and for clarifying what is meant, within the model, by gradual cellular development and memory.

We now explain how the different stages of cellular development — differentiation and senescence — are implemented within this dynamical framework, before describing how the resulting transcriptional changes are quantified using the telegraph model.

3.2 Modelling Differentiation, Senescence and Fate Progression

In the Waddington picture introduced in Sec. 2.1, cell fate corresponds to motion on an effective landscape whose valleys represent stable cellular identities. We now describe how different cell states are implemented in the model by constraining the underlying regulatory degrees of freedom, which may be viewed as "digging" or reshaping valleys in this landscape.

Within this framework, there are two natural and complementary ways to represent differentiation, shown schematically in Figure 6. One can alter the *chromatin fibre* itself, changing which regions of the polymer are accessible for TF binding, or one can modify the composition and diversity of the TF pool interacting with the chromatin. We therefore introduce

two minimal "toy" models of differentiation based on these two ideas, followed by a model of senescence as a qualitatively distinct terminal state. Our aim is not to reproduce any one biological system, but to understand how different architectural changes to the chromatin reshape the effective landscape and thereby influence transcriptional dynamics. The differentiated and senescent states studied here should therefore be interpreted as idealised dynamical regimes rather than direct models of specific cell types. Their purpose is to separate, as cleanly as possible, the roles of chromatin accessibility, TF diversity, and chromatin self-association in shaping transcriptional behaviour.

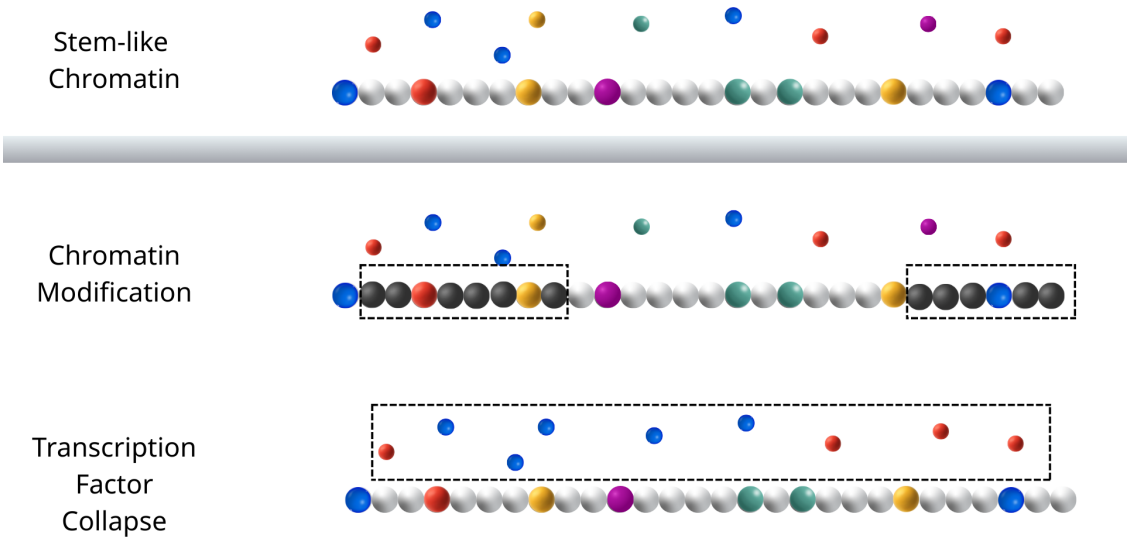

Figure 6. Schematic of the two differentiation protocols introduced in this work. Top: stem-like reference system, consisting of a heterogeneous chromatin fibre with transcription units (TUs; coloured beads) embedded within a background of generic chromatin beads, together with multiple TF species (coloured particles). Middle: A/B polymer differentiation protocol, in which selected chromatin segments are converted into inert type-11 beads (dark beads; dashed boxes). This reduces the availability of binding sites along those regions while leaving the TU positions unchanged. Bottom: TF-collapse differentiation protocol, in which the chromatin fibre is left unchanged but the multicolour TF pool is simplified by collapsing several TF species into fewer remaining classes (dashed box). The two protocols therefore represent distinct idealised routes to differentiation: one acts through chromatin remodelling along the polymer, while the other acts through a reduction in the regulatory degrees of freedom.

Differentiation via chromatin remodelling

One route to differentiation is through chromatin remodelling, in which extended regions of the chromatin become less accessible to proteins involved in transcription, namely TFs. In the landscape picture, this restricts the set of polymer conformations and binding configurations available to the system, thereby narrowing the basin of attraction associated with the stem-like state.

We implement this idea using a simple A/B polymer model loosely inspired by histone modification and DNA methylation in real genomes [56, 57, 58, 59]. The chromatin chain is partitioned into repeating blocks of total length P, each consisting of an A segment of length P_A followed by a B segment of length $P_B = P - P_A$. Beads are assigned to these blocks according to their contour index s along the chain: beads whose indices fall within the first P_A positions of a block belong to the A region, while the remainder belong to the B region.

Beads in A blocks retain their stem-like identities: most remain weak non-specific binding sites (type 12), with a sparse subset designated as transcription units (types 13–17). In contrast, background beads in B blocks are converted to inert beads (type 11), which do not participate in attractive interactions with TFs.

Formally, inert beads are defined by the absence of any attractive interaction, rather than by the absence of all pair forces. Their non-bonded interactions with all other particles remain purely steric and are implemented using the same truncated Lennard–Jones potential as other steric-only pairs, with cut-off

$$r_{i,11}^{(c)} = 2^{1/6}\sigma_{i,11}, \qquad \varepsilon_{i,11} > 0 \quad \text{for all } i, \tag{12}$$

so that type 11 beads retain excluded volume and chain connectivity but possess no attractive tail. They therefore do not bind TFs or contribute to TF-mediated bridging. Importantly, TU beads are preserved unchanged even within B blocks, ensuring that the underlying set of regulatory sites remains fixed while the surrounding chromatin environment is altered.

Physically, this construction creates alternating regions of accessible and inaccessible chromatin. TF-mediated bridging and cluster formation are therefore biased towards A blocks, while B blocks act as passive regions that reduce the effective regulatory freedom of the fibre.

Differentiation via TF colour collapse

A complementary route to differentiation is through changes in TF number and identity rather than in the chromatin itself. In this case, the polymer remains in its stem-like form, but the diversity of TFs capable of binding to it is reduced. In the landscape picture, this lowers the dimensionality of the regulatory degrees of freedom, effectively collapsing multiple shallow valleys into fewer, deeper ones.

We implement this idea through a "TF colour collapse" model. It is motivated by the observed up-regulation of certain TFs, often referred to as "master regulators", that help determine cell fate [60, 61, 62, 63]. In addition, some studies report simultaneous down-regulation of other TF classes, which motivates modelling differentiation here as a reduction in TF diversity rather than simply an increase in TF number for selected colours [64, 65].

The chromatin bead types and their binding affinities are left unchanged, but several TF species are gradually reassigned to a smaller set of remaining species. Concretely, TFs of species 3, 4, and 5 are progressively converted into species 1 and 2, so that the initial five TF colours collapse into two dominant classes. The binding and non-binding states of the

collapsing TFs are reassigned together, preserving the underlying switching dynamics while reducing regulatory diversity.

This procedure increases competition for the same TU classes and suppresses the formation of distinct TF-specific compartments. Clustering and transcriptional activity therefore persist, but with reduced selectivity and a smaller set of stable regulatory patterns.

Senescence via chromatin self-association

Senescence represents a qualitatively distinct terminal state characterised by profound architectural reorganisation, often including the formation of compact heterochromatin foci known as senescence-associated heterochromatin foci (SAHF) [27, 28, 66]. In the landscape picture, this corresponds not merely to a restriction of accessible configurations, but to the introduction of an additional energetic drive towards a deeply collapsed and structurally persistent state.

To model senescence, we again employ an A/B organisation along the polymer, but now modify the physical character of the B blocks. Rather than rendering B-block background beads inert, we convert them into a new bead type (type 18) that is strongly self-attractive. In practice, non-TU beads in B blocks are reassigned from type 12 to type 18, while TU beads (types 13–17) are left unchanged, as in the differentiation model.

These sticky B beads interact with one another through a strong short-ranged Lennard–Jones attraction,

$$\varepsilon_{18,18} = \varepsilon_{\mathrm{SAHF}} = 8.0, \qquad r^{(c)}_{18,18} = r_{\mathrm{bind}}, \tag{13}$$

driving the spontaneous formation of compact B-rich clusters reminiscent of SAHF. We choose $\varepsilon_{\mathrm{SAHF}} = 8.0$ so that the B-block self-attraction is strong enough to induce collapse and stabilise long-lived aggregates on simulation timescales. In the present parameter set, this is taken to be equal in magnitude to the cognate TF–TU binding strength $\varepsilon_{\mathrm{sp}}$, since both are intended to represent strong attractive interactions relative to the weaker non-specific background binding.

To avoid trivial sequestration of TUs into these collapsed B-rich domains, interactions between sticky B beads and TU beads are kept purely steric rather than attractive. Thus,

$$r^{(c)}_{18,c} = 2^{1/6}\sigma_{18,c}, \qquad \varepsilon_{18,c} > 0, \qquad c \in \{13, \ldots, 17\}, \tag{14}$$

so that type 18–TU pairs retain excluded-volume repulsion but possess no attractive tail.

This senescence model therefore differs fundamentally from the differentiation protocols above. Rather than simply reducing accessibility, it introduces an additional energetic drive for chromatin collapse. As a result, the system becomes dominated by long-lived structural aggregates, and transcriptional dynamics are strongly shaped by the imposed architecture.

Gradual fate progression

The models above define distinct target cell states, but differentiation is inherently a dynamical process rather than an instantaneous switch. To capture this, we also implement fate transitions by gradually modifying particle identities during the simulation, so that the system moves continuously from the stem-like interaction rules towards those of the differentiated or senescent state.

In practice, background chromatin beads or TF particles are stochastically reassigned at discrete intervals. For the dynamic A/B differentiation protocol, the polymer is divided into blocks of 400 beads, with the first 200 beads forming the A region and the remaining 200 beads forming the B region. During each recolouring cycle, which occurs every 10^4 timesteps, approximately 5 eligible type-12 beads in the B region are converted to inert type-11 beads, while TUs are left unchanged.

In the TF-collapse protocol, by contrast, polymer bead identities are left unchanged and differentiation is implemented by progressively reducing TF colour diversity in time. At each recolouring cycle, a fraction 0.002 of the remaining TFs belonging to colours 3, 4, and 5 is reassigned into colours 1 and 2. More explicitly, colour-3 TFs are reassigned to colour 1, while colour-4 and colour-5 TFs are reassigned to colour 2, with binding and non-binding states mapped separately as $5 \to 1$, $6 \to 2$, $7 \to 3$, $8 \to 4$, $9 \to 3$, and $10 \to 4$. In this way, the original TF colours 3–5 are progressively absorbed into the first two TF species, so that the TF population becomes less diverse as differentiation proceeds. The chosen recolouring probability and timescale ensure that this change is imposed gradually rather than through an instantaneous quench.

From the landscape perspective, this corresponds to a slow descent down Waddington's landscape into a valley. This gradual implementation is important because it allows the system to pass through a continuous sequence of intermediate states rather than being quenched directly into a final one. In turn, this makes it possible to reverse the differentiation process and examine the extent to which the system can "dedifferentiate", thereby probing forms of cellular memory explored in later sections.

3.3 The Telegraph Model

Having defined the chromatin–TF dynamics and the different cellular-state protocols, we now specify how transcriptional output is extracted from the simulations. To do this, we convert the microscopic dynamics into binary activity traces for individual transcription units (TUs) and analyse these traces using a simple two-state telegraph model [22, 23, 24]. The central object is a time series $s_i(t) \in \{0, 1\}$ for each TU bead i, where $s_i(t) = 1$ denotes that TU i is transcriptionally active at time t, and $s_i(t) = 0$ denotes inactivity. Here, time is measured in units of saved trajectory frames rather than in the finer integration timestep or the intermediate TF switching interval. This provides a coarse-grained transcriptional readout of the underlying 3D chromatin–TF dynamics. We first define how the TU activity

traces are constructed from the simulation trajectories, then infer telegraph parameters for simulations in which the chromatin state is held fixed, and finally extend the same framework to dynamic simulations in which the chromatin state evolves over time.

Defining TU activity

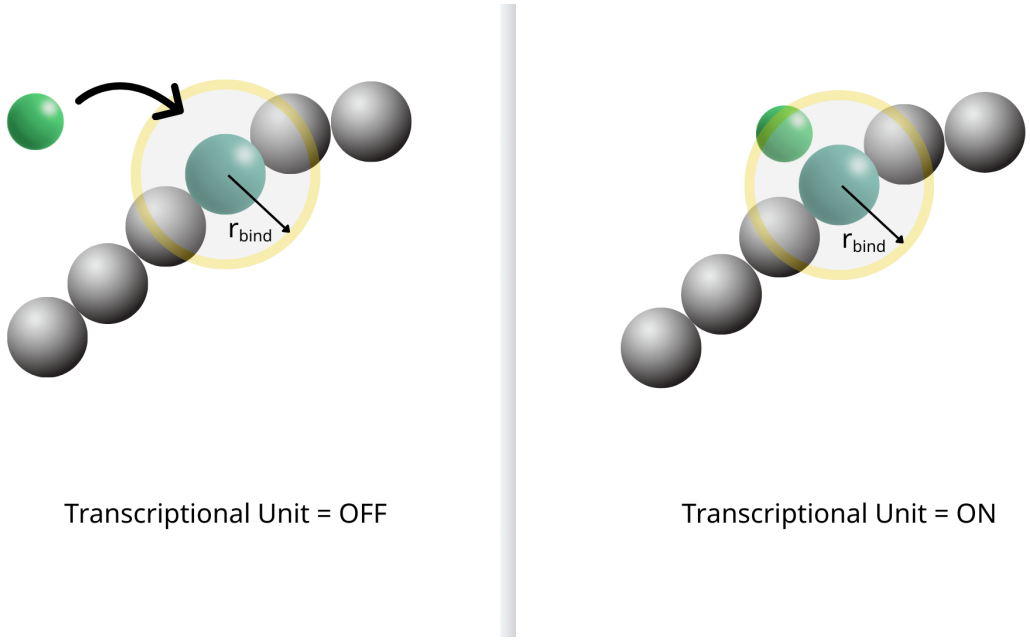

Figure 7. Schematic representation of the two activity states in which a TU can exist at any given time t. If the TU is ON, $s_i(t) = 1$. If the TU is OFF, $s_i(t) = 0$. Over the course of the simulation, each TU therefore has an associated binary activity trace $s_i(t)$.

We define TU activity operationally using a binding criterion between a TU bead and its cognate binding TFs. For a given TU bead i of type $(12+c)$, we declare it active at time t if at least one cognate binding TF of type $(2c-1)$ lies within the binding range r_{bind} (Figure 7),

$$s_i(t) = \begin{cases} 1, & \exists\, p \in \mathcal{P}_{2c-1} \text{ s.t. } |\mathbf{r}_p(t) - \mathbf{r}_i(t)|_{\text{PBC}} < r_{\text{bind}}, \\ 0, & \text{otherwise}, \end{cases} \tag{15}$$

where $\mathcal{P}_{2c-1}$ denotes the set of TF particles of the cognate binding type and $|\cdot|_{\text{PBC}}$ indicates that distances are computed with periodic boundary conditions.

This definition is chosen to align directly with the microscopic interaction rules of the model: a TU is counted as active precisely when a cognate binding TF is present within the same range that mediates attractive TF–chromatin interactions (Sec. 3.1). In this sense, TU activity is not introduced as an additional dynamical variable, but is inferred directly from the underlying chromatin–TF dynamics. We adopt the minimal criterion of at least one cognate binding TF within r_{bind}, rather than imposing a stronger threshold based on TF number or cluster membership, so as to avoid introducing further arbitrary definitions beyond those already built into the interaction model. The raw transcriptional output recorded at each frame is therefore $(t,\ i,\ s_i(t),\ \text{type}_i)$.

Inferring telegraph rates in stationary simulations

Given a binary activity trace $s_i(t)$, the telegraph model represents TU activity as a two-state continuous-time Markov process,

$$\text{OFF} \underset{k_{\text{off}}}{\overset{k_{\text{on}}}{\rightleftharpoons}} \text{ON}.$$

In this model, OFF→ON and ON→OFF transitions occur as Poisson processes with constant rates k_{on} and k_{off}. In practice, however, the simulation output is sampled at discrete saved frames separated by a sampling interval Δt, which we take to be one frame. We therefore first estimate transition probabilities per frame and then convert them to continuous-time rates.

For a given TU trace $\{s_t\}_{t=0}^{F-1}$ over F frames, we count the number of adjacent-frame transitions in which the activity switches from $0 \to 0$, $0 \to 1$, $1 \to 0$, or $1 \to 1$:

$$\begin{aligned} N_{00} &= \sum_t \delta_{s_t,0}\,\delta_{s_{t+1},0}, \quad N_{01} = \sum_t \delta_{s_t,0}\,\delta_{s_{t+1},1}, \\ N_{10} &= \sum_t \delta_{s_t,1}\,\delta_{s_{t+1},0}, \quad N_{11} = \sum_t \delta_{s_t,1}\,\delta_{s_{t+1},1}. \end{aligned} \tag{16}$$

These satisfy $N_{00} + N_{01} + N_{10} + N_{11} = F - 1$. From these counts, we estimate the per-frame transition probabilities

$$P_{01} = \Pr(s_{t+1} = 1 \mid s_t = 0) = \frac{N_{01}}{N_{00} + N_{01}}, \qquad P_{10} = \Pr(s_{t+1} = 0 \mid s_t = 1) = \frac{N_{10}}{N_{10} + N_{11}}, \tag{17}$$

with the understanding that these expressions are undefined if the denominators vanish, for example if a TU is never observed in the corresponding state over the analysed interval.

To map these discrete probabilities onto continuous-time rates, we use the standard relation for a Poisson process: the probability of *no* event in a time interval Δt is $\exp(-k\Delta t)$. Identifying $1 - P$ with this no-event probability yields

$$k_{\text{on}} = -\frac{1}{\Delta t}\ln(1 - P_{01})\,, \qquad k_{\text{off}} = -\frac{1}{\Delta t}\ln(1 - P_{10})\,, \tag{18}$$

from which we define the corresponding telegraph timescales,

$$\tau_{\text{on}} = \frac{1}{k_{\text{off}}}, \qquad \tau_{\text{off}} = \frac{1}{k_{\text{on}}}, \qquad \tau_{\text{corr}} = \frac{1}{k_{\text{on}} + k_{\text{off}}}. \tag{19}$$

These derived timescales are later used to characterise transcriptional persistence and the dynamical signatures of memory.

In addition to the switching rates, we quantify transcriptional output using the empirical mean ON fraction of each TU,

$$p_{\text{on}} = \frac{1}{F}\sum_{t=0}^{F-1} s_t, \tag{20}$$

where F is the number of sampled frames. For a two-state telegraph process in steady state, this quantity is predicted to satisfy

$$p_{\mathrm{on}}^{(\mathrm{tele})} = \frac{k_{\mathrm{on}}}{k_{\mathrm{on}} + k_{\mathrm{off}}}, \tag{21}$$

providing a simple closure relation. Agreement between the measured p_{on} and the value inferred from the rates indicates that the binary activity traces are reasonably well described by a two-state process on the sampled timescales, whereas systematic deviations suggest that the telegraph model is missing some features of the underlying dynamics. In later sections, we therefore use this closure relation as an internal consistency check before interpreting the inferred switching rates and correlation times physically.

Telegraph inference for dynamic simulations

The above telegraph parameters are appropriate for approximately stationary simulations in which the chromatin state is held fixed throughout the run. For the dynamic differentiation and senescence protocols introduced above, however, the chromatin state evolves over time, so the transcriptional kinetics must also be inferred in a time-resolved manner. To quantify this progression, we extend the telegraph analysis using sliding windows. This windowed approach balances two competing requirements: each window must contain enough switching events for the inferred rates to remain meaningful, while still being short enough to resolve the gradual evolution of the transcriptional state during differentiation and senescence.

Concretely, we choose a window size W (in frames) and a stride S (in frames), and define windows $\{t = t_0, \ldots, t_0 + W - 1\}$ with $t_0 = 0, S, 2S, \ldots$. Within each window, we compute the mean activity $p_{\mathrm{on}}(t_0)$ from Eq. (20), the transition counts $N_{ab}(t_0)$ from Eq. (16), the corresponding probabilities $P_{01}(t_0)$ and $P_{10}(t_0)$ from Eq. (17), and the time-dependent rates $k_{\mathrm{on}}(t_0)$ and $k_{\mathrm{off}}(t_0)$ via Eq. (18). We assign each window a time coordinate equal to its midpoint,

$$t_{\mathrm{center}} = t_0 + \frac{W - 1}{2}, \tag{22}$$

and report kinetic trajectories $k_{\mathrm{on}}(t_{\mathrm{center}})$, $k_{\mathrm{off}}(t_{\mathrm{center}})$, $p_{\mathrm{on}}(t_{\mathrm{center}})$, and $\tau_{\mathrm{corr}}(t_{\mathrm{center}})$ over the course of cellular development.

This time-resolved extension allows the same telegraph framework to be applied consistently to both static and evolving chromatin states. It therefore provides a common language for comparing stem-like, differentiated, senescent, and dynamically progressing simulations, and forms the basis for the later analysis of activity redistribution and cellular memory.

4 Results and Discussion

Having established the simulation framework and transcriptional readout, we now turn to the central question of this project: how does chromatin organisation reshape transcriptional dynamics, and how do those dynamics relate to stages of cellular development on Waddington's landscape?

We address this in three steps. First, we compare fixed chromatin states — stem-like, differentiated, and senescent — in order to establish the phenomenon of fossilisation and determine how transcriptional activity is redistributed across TUs in each case. Second, we follow differentiation dynamically, allowing the system to "roll down" the landscape in time, and quantify how the transcriptional programme departs from the stem-like state, how activity becomes redistributed across TUs, and how commitment emerges. Finally, we investigate cellular memory by reversing differentiation and asking whether the differentiated transcriptional state is undone or retained, and how this depends on transcription-factor abundance.

4.1 Rising inequality associated with differentiation and senescence

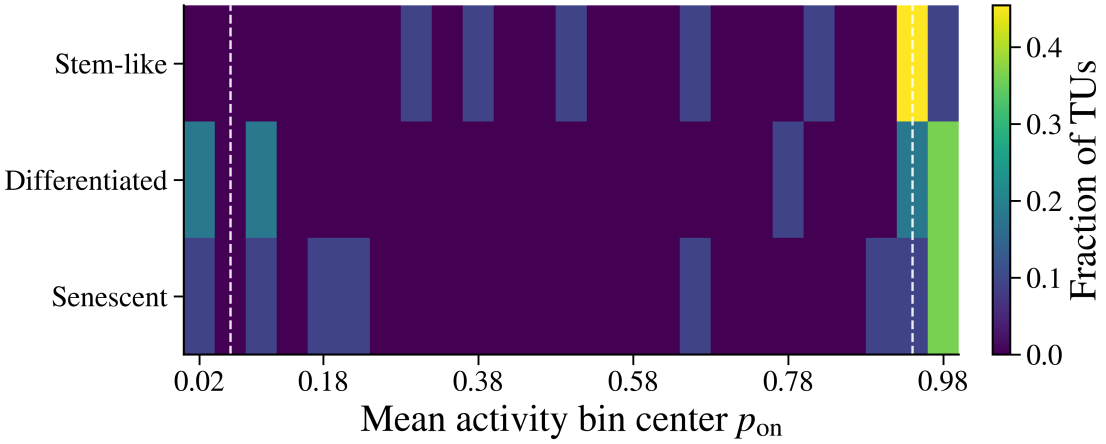

Figure 8. Binned density heatmap of mean activity p_{on} across TUs for each chromatin state. Colour indicates the fraction of TUs in each bin. Stem-like chromatin shows a broad distribution across intermediate values, whereas differentiated and senescent states display depletion of intermediate bins and accumulation near extreme ON or OFF states. Dashed vertical lines indicate the locking thresholds.

We begin by analysing simulations in which the form of chromatin organisation remains fixed throughout the entire run. By "static" we mean that the system is initialised in a given chromatin state — stem-like, differentiated, or senescent — and then evolved in that state without further modification. In this section, we hold the TF count fixed at 40 TFs per type. This value was chosen to place the system in a regime where TF-mediated bridging and clustering remain robust, while still allowing competition between TUs for a limited TF pool. We also focus on the chromatin modification differentiation protocol, since the corresponding static behaviour is qualitatively similar when TF collapse is used, which we return to in Sec. 4.2.

Our aim here is not simply to compare mean transcriptional output between conditions, which in many cases changes only modestly. Instead, we ask how activity is redistributed

across TUs at each stage of cellular development. We define *fossilisation* as a redistribution of TU activity away from intermediate, flexible switching states and towards more extreme, effectively locked configurations. We first show this directly in the distribution of TU activities, then quantify the increase in locked fractions, next examine how activity becomes concentrated into a restricted subset of TUs, and finally relate this redistribution to the underlying burst dynamics.

Figure 8 shows the distribution of the mean activity p_{on} across TUs in each chromatin state. In the stem-like state, activity spans a broad range of intermediate values, forming a relatively continuous distribution, although there is still a substantial fraction of highly active TUs. There are also very few low-activity TUs, consistent with the idea that most TUs remain able to become active at some point. In contrast, both differentiated and senescent chromatin show a clear depletion of intermediate bins and a corresponding accumulation near the extremes. Many TUs are either almost always inactive or almost always active, with fewer occupying intermediate switching regimes.

This polarisation is the most direct static signature of fossilisation: flexible intermediate states become less populated once chromatin differentiates or enters senescence. To quantify how large this shift towards extreme states is, we next classify TUs according to simple locking thresholds.

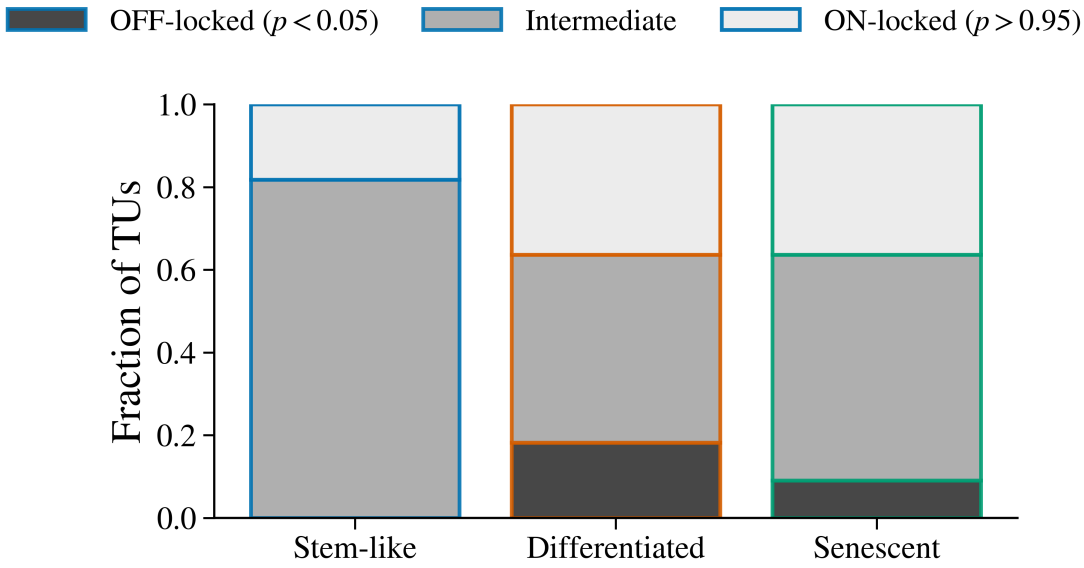

Figure 9. Fraction of TUs classified as OFF-locked ($p_{\mathrm{on}} < 0.05$), intermediate, or ON-locked ($p_{\mathrm{on}} > 0.95$) in each chromatin state. Differentiated and senescent states show an increased proportion of TUs in extreme states compared to the stem-like condition.

Figure 9 quantifies this effect by classifying TUs as OFF-locked ($p_{\mathrm{on}} < 0.05$), ON-locked ($p_{\mathrm{on}} > 0.95$), or intermediate. In the stem-like state, the majority of TUs lie in the intermediate band, indicating that they are neither permanently silent nor permanently active. In differentiated and senescent states, a substantially larger fraction of TUs are locked, particularly in the OFF category, and the intermediate population is correspondingly reduced.

Together, Figures 8 and 9 show that differentiation and senescence shift the system from a more flexible configuration to one dominated by extreme states. The next question is how this redistribution is organised across the TU population: is activity suppressed gradually, or does it become concentrated into a restricted dominant subset?

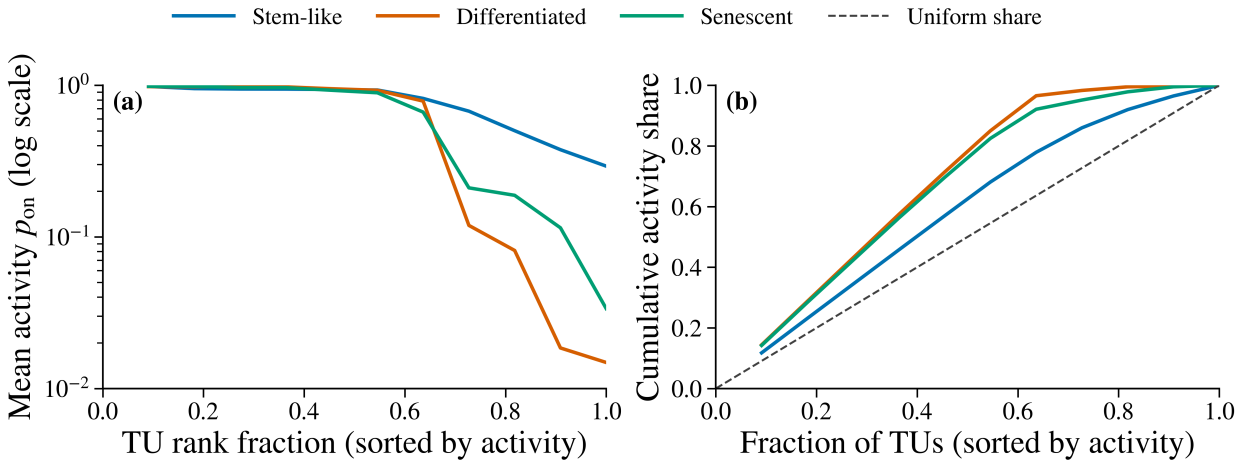

Figure 10. Static activity redistribution across TUs in stem-like, differentiated, and senescent chromatin states. (a) Rank–activity curve showing mean activity p_{on} as a function of TU rank fraction (sorted in descending order). The logarithmic scale emphasises differences in the low-activity tail. Differentiated and senescent states show a sharper collapse of activity in the lower-ranked TUs compared to the stem-like state. (b) Cumulative activity share across TUs, where TUs are sorted by mean activity p_{on} and the cumulative fraction of total activity is plotted against the fraction of TUs. The dashed line indicates a uniform distribution in which all TUs contribute equally. Differentiated and senescent states exhibit a steeper rise, indicating concentration of activity into a smaller subset of TUs.

Figure 10(a) addresses this by ranking TUs in descending order of their mean activity and plotting p_{on} against rank fraction on a logarithmic scale. All three conditions exhibit a highly active top subset. However, the behaviour of the lower-ranked TUs differs markedly. In the stem-like state, activity decays gradually, and even lower-ranked TUs retain moderate transcriptional activity. In differentiated and senescent states, by contrast, activity collapses sharply beyond a certain rank fraction, with many TUs exhibiting orders-of-magnitude lower p_{on}. This reveals a much stronger hierarchy in TU activity lower down the landscape.

Figure 10(b) shows the cumulative fraction of total activity carried by the top fraction of ranked TUs. The dashed diagonal corresponds to a uniform distribution in which all TUs contribute equally. The stem-like curve lies closest to this diagonal, indicating that activity is relatively broadly distributed across the TU population. In contrast, both differentiated and senescent states rise more steeply, meaning that a smaller subset of TUs accounts for a disproportionately large fraction of total transcriptional output. In this sense, the stem-like state is comparatively egalitarian, whereas differentiation and senescence concentrate activity into a smaller subset of dominant TUs.

Taken together, Figures 8–10 show that fossilisation has two connected aspects. First, intermediate switching regimes are depleted and extreme ON/OFF states become more common.

Second, the remaining transcriptional activity becomes increasingly concentrated into a restricted subset of TUs. We now ask what dynamical change underlies this redistribution. Because the closure relation (Eq. 21) is satisfied well across the static runs, the binary TU activity traces are reasonably well described by an effective two-state process on the sampled timescales, so the telegraph parameters provide a meaningful summary of the switching dynamics.

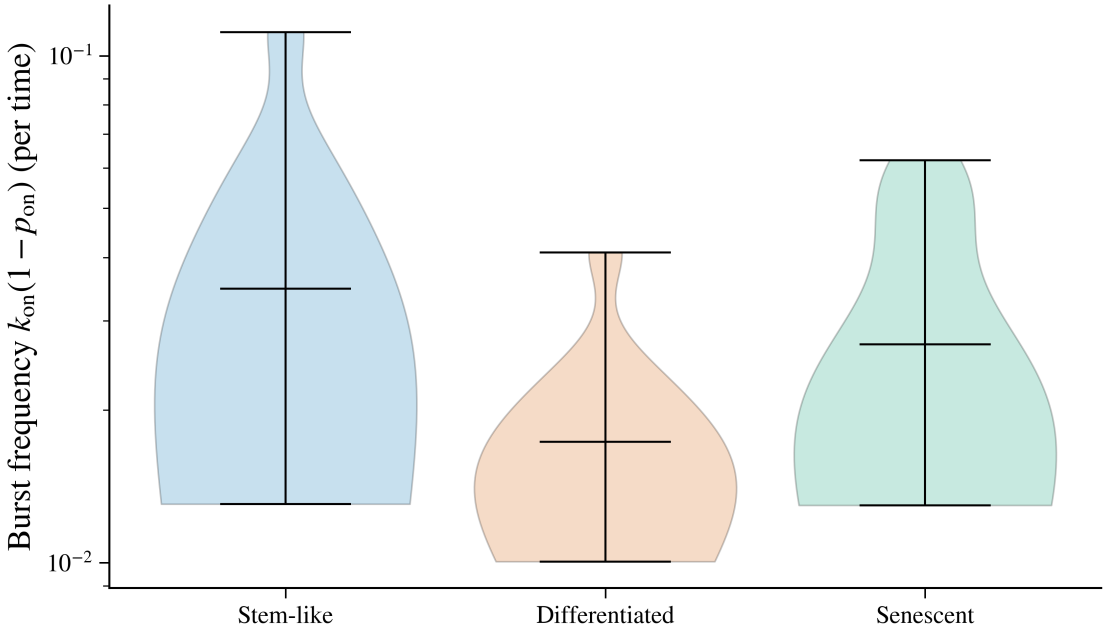

Figure 11. Distribution of burst initiation frequencies $k_{on}(1 - p_{on})$ across TUs for each chromatin state (logarithmic scale). Stem-like chromatin shows higher and broader initiation frequencies, while differentiated and senescent states display suppression of burst initiation.

Figure 11 shows the distribution of burst initiation frequencies, defined from the telegraph model as $k_{on}(1 - p_{on})$. This quantity measures how often a TU initiates transcriptional episodes. The stem-like state exhibits the highest median burst frequency and a broad distribution, consistent with frequent transitions into the ON state. In differentiated and senescent states, the distribution shifts downwards, with many TUs exhibiting very low initiation frequencies.

Lower burst frequency implies that TUs activate less often and are therefore more likely to remain in stable OFF configurations for extended periods. This provides a dynamical interpretation of the redistribution seen above: intermediate switching regimes disappear primarily because activation itself is suppressed. The static analysis therefore identifies three linked effects of altered chromatin organisation. Intermediate activity states are depleted, transcription becomes concentrated into a smaller subset of dominant TUs, and burst initiation is suppressed, stabilising inactive configurations. This is precisely the pattern we refer to as fossilisation.

A notable and somewhat surprising feature of these results is that the senescent state does not always exhibit a stronger degree of fossilisation or inequality than the differentiated state,

despite the strong self-association of B blocks in the senescence model. At first sight this may seem counterintuitive, since SAHF-like compaction might be expected to produce the most strongly locked transcriptional landscape. A plausible interpretation is that the two models constrain transcription in qualitatively different ways. In the differentiated A/B model, B blocks are rendered inert, so large regions of the polymer lose their ability to participate in productive TF binding, which may channel activity towards the remaining accessible regions and thereby enhance inequality. In the senescence model, by contrast, B blocks become strongly self-attractive, driving large-scale chromatin collapse and reorganisation, but not necessarily imposing the same selective redistribution of TF accessibility. The dominant effect may therefore be architectural aggregation rather than a clean pruning of the accessible regulatory space. Thus, stronger chromatin compaction does not automatically imply stronger transcriptional inequality. Indeed, the collapse of chromatin that occurs during senescence may in some cases help bring separated but binding-competent regions closer to each other, increasing the chances of activity initiation. This is a subtle but important distinction that we cannot pursue in greater detail in the present work.

Within the literature we have considered, differentiation has not previously been described explicitly in terms of a redistribution of transcriptional activity across TUs in the sense of the inequality developed here. A substantial body of single-cell work has shown that stem-like states exhibit elevated transcriptional variability that becomes progressively constrained during commitment to a particular identity [67, 68, 69]. These studies, however, typically quantify variability across cells or along developmental trajectories, rather than asking how activity is partitioned across TUs within a fixed chromatin state. In other words, reduced TU variability has been well documented, but not recast as a measurable concentration of transcriptional output into a restricted subset of TU clusters. Our notion of fossilisation therefore provides a complementary perspective: differentiation is not merely associated with reduced variability, but with a structural redistribution of transcriptional dynamics away from intermediate switching regimes and towards effectively locked ON/OFF configurations. Framing commitment in terms of transcriptional-unit inequality therefore extends the Waddington-landscape description of cell fate in a new direction.

These static results, however, describe only the endpoints of each chromatin state. They do not reveal how the system moves from the stem-like state to a differentiated one. To address this, we now follow differentiation dynamically and analyse how transcriptional parameters evolve as the system rolls down Waddington's landscape.

4.2 Rolling down Waddington's landscape in two ways

In the previous section, we compared static chromatin states and showed that differentiation reshapes transcription by redistributing activity across TUs, thereby increasing inequality in TU activity. We now move from comparing fixed endpoints to analysing the differentiation process itself. Instead of initialising directly in a differentiated configuration, we begin in the stem-like state and gradually differentiate according to either the chromatin-remodelling protocol or the TF colour-collapse protocol described in Sec. 3.2. As before, we keep the

initial TF pool fixed at 40 TFs per type, so that the only changing ingredient is the differentiation mechanism itself. We also restrict attention here to differentiation; senescence represents a terminal regime beyond the landscape descent that we wish to study.

Our aim is to understand how the transcriptional programme evolves as the system "rolls down" Waddington's landscape. We proceed in four steps. First, we quantify how far the transcriptional programme moves from the initial stem-like state. Second, we ask whether activity becomes more unevenly distributed across TUs within each run. Third, we examine whether different runs become more or less variable as differentiation proceeds. Finally, we connect these changes to the underlying switching dynamics and to the emergence of commitment.

To quantify programme-level change, we track the full pattern of TU activities in time. In each time window, we compute the window-averaged activity of every TU i,

$$p_i(t) = \frac{1}{W} \sum_{\tau \in \text{window}(t)} s_i(\tau), \tag{23}$$

and collect these into a programme vector

$$\mathbf{p}(t) = \big(p_1(t), p_2(t), \ldots, p_N(t)\big). \tag{24}$$

This vector represents the instantaneous transcriptional identity of the system. To quantify how far the system has moved from its stem-like starting point $\mathbf{p}_0$, we compute the cosine distance

$$\mathrm{D}(t) = 1 - A(t) = 1 - \cos\big(\mathbf{p}(t),\, \mathbf{p}_0\big) = 1 - \frac{\mathbf{p}(t) \cdot \mathbf{p}_0}{\|\mathbf{p}(t)\|\, \|\mathbf{p}_0\|}. \tag{25}$$

We refer to $D(t)$ as the *programme drift*. It is zero at the start of the trajectory and increases as the transcriptional pattern departs from the stem-like configuration. As in the static analysis, we checked that the closure relation is satisfied well across the dynamic runs, indicating that the telegraph parameters continue to provide a useful coarse-grained summary of the switching dynamics.

Figure 12 shows that both differentiation protocols produce a steady increase in programme drift. The system does not merely fluctuate locally around its initial state; it progressively departs from its original stem-like transcriptional identity. TF collapse generates a larger overall drift, indicating a stronger global reorganisation of which TUs are active, whereas chromatin remodelling produces a more moderate but still clear departure. This establishes that differentiation changes the programme, but not yet how that change is structured across the TU population.

To address this, we next examine how uneven TU activity becomes within each simulation. We define the *heterogeneity* as the variance of $p_i(t)$ across TUs, normalised by the Bernoulli variance associated with the global mean activity in that window,

$$p_{\text{on}}(t) = \langle p_i(t) \rangle_i, \tag{26}$$

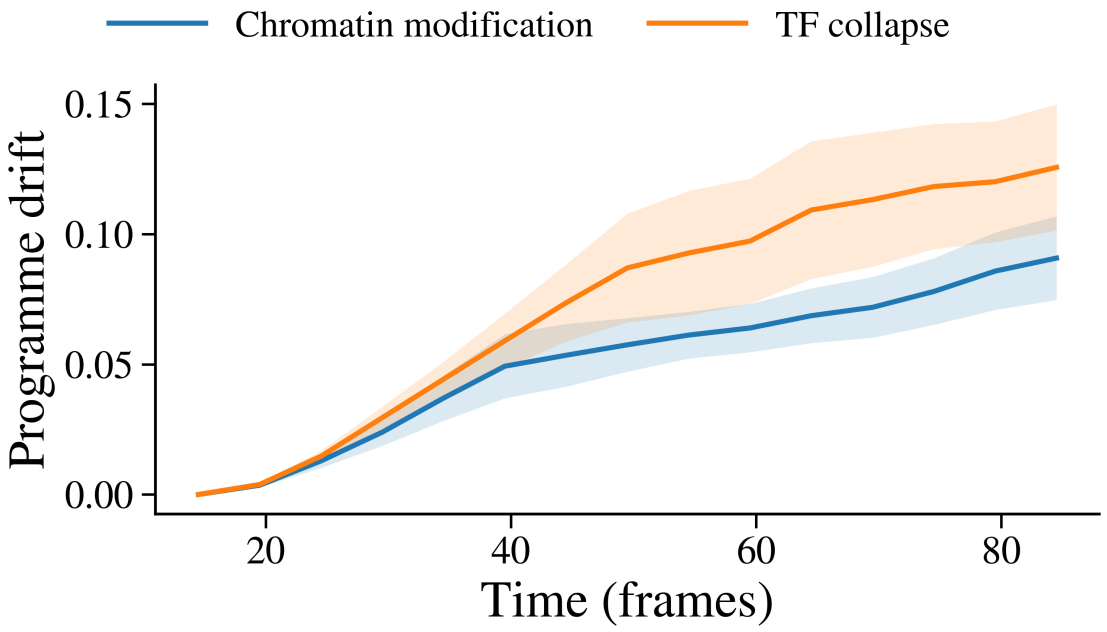

Figure 12. Programme drift during dynamic differentiation for the chromatin-remodelling protocol and the TF-collapse protocol (mean ± SEM across runs). Drift is defined as a cosine distance from the initial stem-like transcriptional programme. Both protocols show progressive departure from the initial state, with TF collapse producing a larger overall drift.

$$\text{Heterogeneity}(t) = \frac{\text{Var}_i[p_i(t)]}{p_{\text{on}}(t)\big(1 - p_{\text{on}}(t)\big)}. \tag{27}$$

This normalisation ensures that heterogeneity captures redistribution across TUs rather than trivial changes in overall activity. Physically, heterogeneity measures how unequal TU activities are within a given realisation. This is distinct from programme drift: drift measures how far the overall transcriptional pattern moves from the initial stem-like state, whereas heterogeneity measures how unevenly activity is distributed across TUs within a given state.

As shown in Fig. 13(a), heterogeneity rises strongly during differentiation for both protocols. Activity becomes increasingly polarised: some TUs dominate while others are suppressed. The accompanying interquartile range across TUs, shown in Fig. 13(b), confirms that this broadening of the activity distribution is not driven by rare extremes but reflects a genuine increase in TU-activity inequality. This is the dynamic counterpart of the fossilisation established in Sec. 4.1: as differentiation proceeds, TUs are progressively separated into low-activity and high-activity regimes. Chromatin remodelling ultimately produces slightly higher heterogeneity, consistent with a more stratified final state.

This also clarifies why chromatin remodelling can produce a more unequal differentiated state without producing the largest programme drift. The two quantities capture different features of the dynamics. Chromatin remodelling largely preserves which TUs tend to be active, but amplifies the differences between them, pushing some further ON and others further OFF. TF collapse, by contrast, reshapes the regulatory balance more globally, leading to a larger

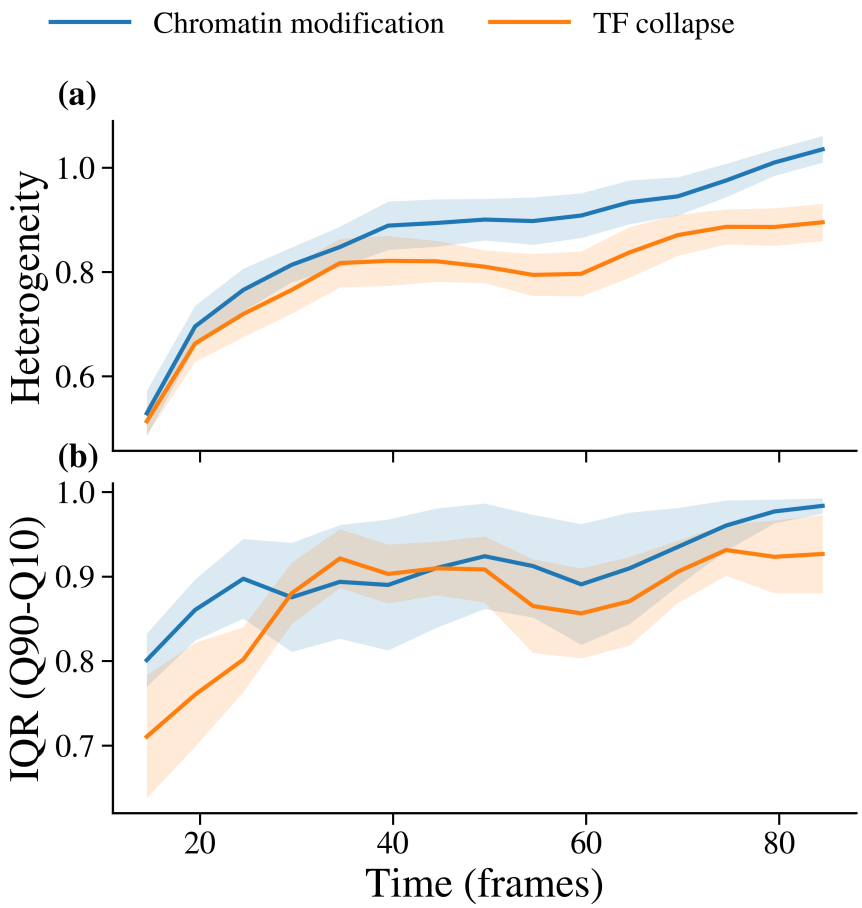

Figure 13. Evolution of within-run activity spread during differentiation (mean ± SEM across runs). (a) Heterogeneity, defined as the normalised variance across TUs, increases with time for both protocols. (b) The interquartile range across TUs shows the same broadening trend, confirming increasing polarisation of activity.

reorganisation of which TUs dominate. In this sense, chromatin remodelling deepens existing inequalities, whereas TF collapse reshuffles them.

Heterogeneity measures variability *within* a single simulation. Plasticity, by contrast, measures variability *across* simulations. To quantify plasticity, we examine how much the window-averaged activity of each TU varies between independent realisations of the same protocol. For each TU, we compute the cross-simulation variance and normalise it,

$$\text{Plasticity}(t) = \left\langle \frac{\text{Var}_r[p_{i,r}(t)]}{\mu_i(t)\big(1-\mu_i(t)\big)} \right\rangle_i , \tag{28}$$

where $\mu_i(t)$ is the mean activity of TU i across runs. While heterogeneity asks how unequal TUs are within one realisation, plasticity asks whether different realisations settle into different active subsets of TUs.

Figure 14(a) shows that plasticity is not simply monotonic during differentiation. Here, however, "plasticity" should be interpreted mainly as a measure of cross-run variability in

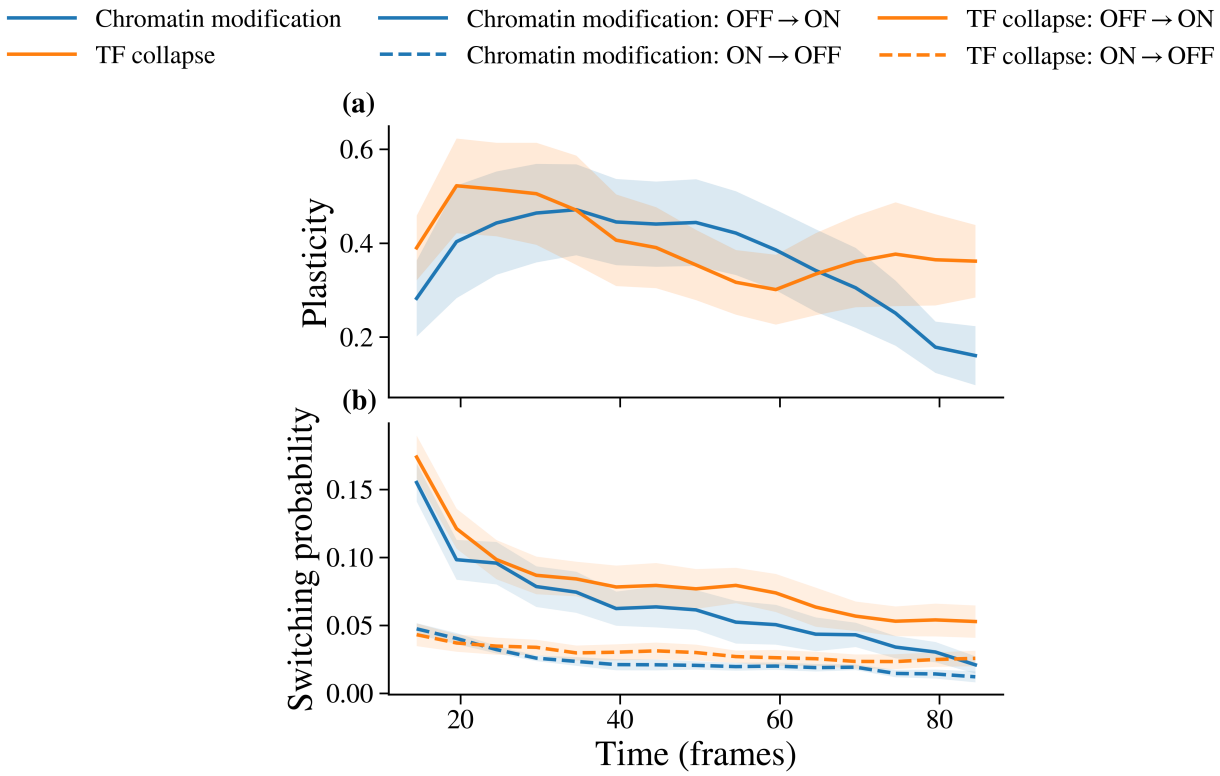

Figure 14. Plasticity and switching dynamics during differentiation. (a) Cross-run plasticity shows an initial rise followed by a later decline, consistent with an early exploratory phase before increasing commitment at later times. (b) OFF→ON transition probabilities decrease strongly, whereas changes in ON→OFF probabilities are weaker, indicating that the loss of plasticity is driven primarily by suppression of reactivation rather than by enhanced deactivation.

the transcriptional programmes selected by the system, rather than as a literal measure of biological cell-fate plasticity. Because independent runs begin from different stochastic initial conditions and random seeds, part of this variance reflects the stochasticity of initial conditions within the model. With that caveat in mind, the early rise in plasticity suggests an exploratory phase in which different realisations stabilise different active subsets of TUs, whereas the later decline indicates increasing commitment to a smaller and more reproducible set of configurations.

The extent of this decline depends on the differentiation mechanism. Chromatin remodelling produces a much stronger late-time reduction in plasticity, consistent with its higher heterogeneity and more stratified TU activity. By restricting chromatin accessibility, this protocol limits the number of viable transcriptional configurations and thereby reduces cross-simulation variability. In contrast, TF collapse retains a larger residual plasticity at late times. Although it substantially reshapes the transcriptional programme and clearly drifts away from the stem-like state, it does not strongly restrict which TUs can in principle remain active. All regions of chromatin remain accessible, and multiple distinct active subsets are still compatible with the modified TF pool. As a result, different runs can stabilise different

configurations even at late times. The decrease in plasticity should therefore be interpreted not as proof of reduced biological plasticity in a strict sense, but as evidence that differentiation reduces the cross-run degeneracy of accessible transcriptional programmes, making the final state more reproducible despite the stochasticity of initialisation and cluster formation.

To understand what drives this loss of plasticity, we now examine the underlying switching dynamics. Within each window, we measure the probabilities of OFF→ON and ON→OFF transitions. As shown in Fig. 14(b), differentiation is accompanied primarily by a strong suppression of OFF→ON transitions, while ON→OFF changes more modestly. Physically, this means that once a TU becomes inactive, it is increasingly unlikely to reactivate. Plasticity is therefore lost not because active TUs are aggressively shut down, but because silent TUs are increasingly unable to reactivate.

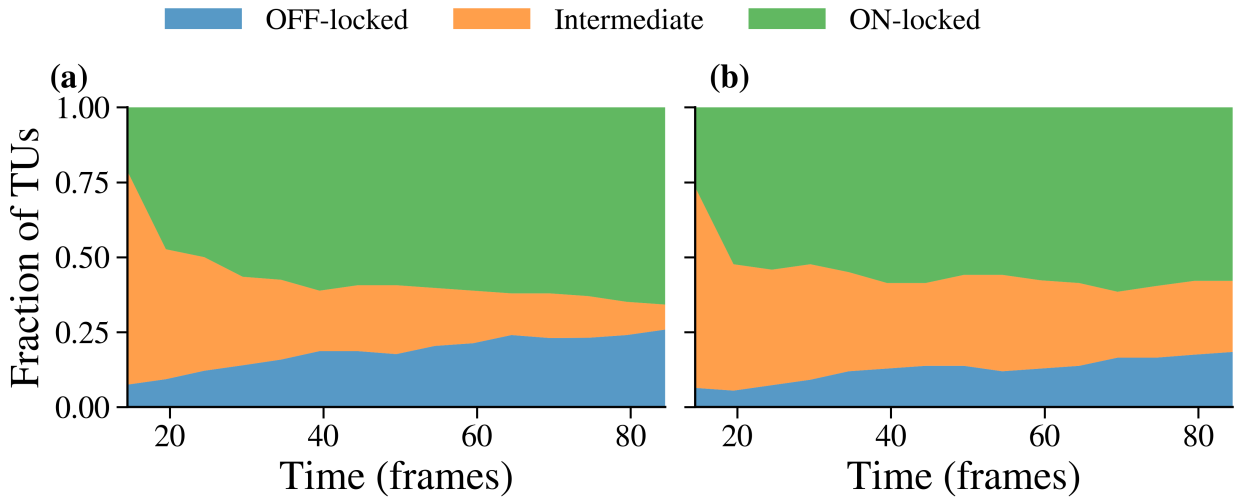

Figure 15. Commitment curves during dynamic differentiation. The fraction of TUs in OFF-locked, intermediate, and ON-locked regimes is shown as a function of time. Differentiation progressively depletes the intermediate population and increases the locked fractions, particularly OFF-locked TUs.

Finally, we visualise commitment directly by classifying TUs according to whether their window-averaged activity is close to zero (OFF-locked), close to one (ON-locked), or intermediate. Figure 15 shows that during differentiation the intermediate population steadily shrinks, while both locked fractions increase. The switching dynamics explain this behaviour naturally. Early in the trajectory, OFF→ON transitions are still sufficiently frequent to allow a subset of TUs to become strongly active, which contributes to the early rise in the ON-locked fraction at the expense of the intermediate population. As differentiation proceeds, OFF→ON transitions are strongly suppressed, so that TUs that fall inactive are increasingly unlikely to reactivate. In this later regime, silent states accumulate more persistently, producing the continued rise in the OFF-locked fraction. Commitment therefore emerges through two stages: early selection and stabilisation of a restricted active subset, followed by progressive persistence of silent states as reactivation becomes increasingly rare. The depletion of intermediate switching regimes reflects the overall reduction in dynamical exploration as the system settles into a differentiated configuration.

Taken together, Figs. 12–15 show a coherent physical picture. Differentiation moves the system away from its stem-like transcriptional identity (programme drift), increases the inequality of activity across TUs (heterogeneity and IQR), reduces variability across independent realisations (plasticity), and progressively locks TUs into extreme regimes (commitment curves). The dominant kinetic change is suppression of OFF→ON transitions, implying that landscape descent commits the system primarily by raising the barrier to reactivation rather than by enforcing an abrupt shutdown of active TUs.

This picture mirrors what is observed in many biological systems. Differentiation is often associated not with an instantaneous global shutdown of gene activity, but with a progressive restriction of regulatory possibilities [70]. Single-cell studies further show that exit from stem-like states and subsequent differentiation occur through staged transitions in which transcriptional variability is reshaped before stabilising into distinct lineage-specific programmes [67, 68]. As a cell commits to a lineage, alternative transcriptional programmes become increasingly difficult to access, and transcriptional state becomes more predictive of eventual fate [70]. Yamanaka's classic reprogramming study likewise shows that differentiation involves progressive silencing of pluripotency-associated genes alongside increasing chromatin constraint, reflecting a gradual tightening of the regulatory landscape and reduced accessibility of alternative transcriptional states [71]. In our model, this appears as a suppression of OFF→ON transitions. The differentiated state therefore emerges not from instantaneous shutdown of activity, but from a gradual loss of plasticity.

This naturally leads to the next question: once a differentiated transcriptional state has formed, how strongly is it remembered? To answer this, we now reverse differentiation and investigate whether the system returns towards the stem-like state or remains trapped near the differentiated configuration.

4.3 Canalisation: Understanding Cellular Memory

In the previous sections, we showed that differentiation reshapes transcriptional dynamics in two connected ways: static analysis revealed fossilisation into extreme ON/OFF regimes (Sec. 4.1), while dynamic analysis showed progressive commitment driven primarily by suppression of OFF→ON transitions (Sec. 4.2). Across that discussion, varying TF number produced only quantitative changes in heterogeneity and plasticity; the qualitative picture of differentiation remained intact. Where TF number becomes more decisive is in *memory*: how strongly a differentiated transcriptional state resists being undone.

To investigate this, we implement a hysteresis protocol. Starting from a stem-like state, we differentiate using the TF-collapse protocol (Sec. 3.2) and then gradually reverse the differentiation rules. We then ask whether the transcriptional identity returns towards the original stem-like configuration or remains trapped near the differentiated state. In the language of Waddington's landscape (Fig. 1), weak memory corresponds to a shallow differentiated valley that can be escaped upon reversal, whereas strong memory corresponds to a deep, "canalised" basin that retains its identity even when the external driver of differentiation is undone.

We focus exclusively on TF-collapse differentiation. This ensures that TF abundance is the primary control parameter shaping the effective landscape. Chromatin accessibility itself is left unchanged, so any change in reversibility can be attributed directly to TF number. In this way, the hysteresis protocol isolates one control axis of the model, avoiding the additional ambiguity that would arise if chromatin accessibility and TF composition were varied simultaneously during reversal.

We begin with the telegraph correlation time τ_{corr}, defined previously in Eq. (19),

$$\tau_{\mathrm{corr}} = \frac{1}{k_{\mathrm{on}} + k_{\mathrm{off}}},$$

where the effective switching rates k_{on} and k_{off} are inferred using Eqs. (17)–(18). Physically, τ_{corr} measures how long TU activity states persist before decorrelating.

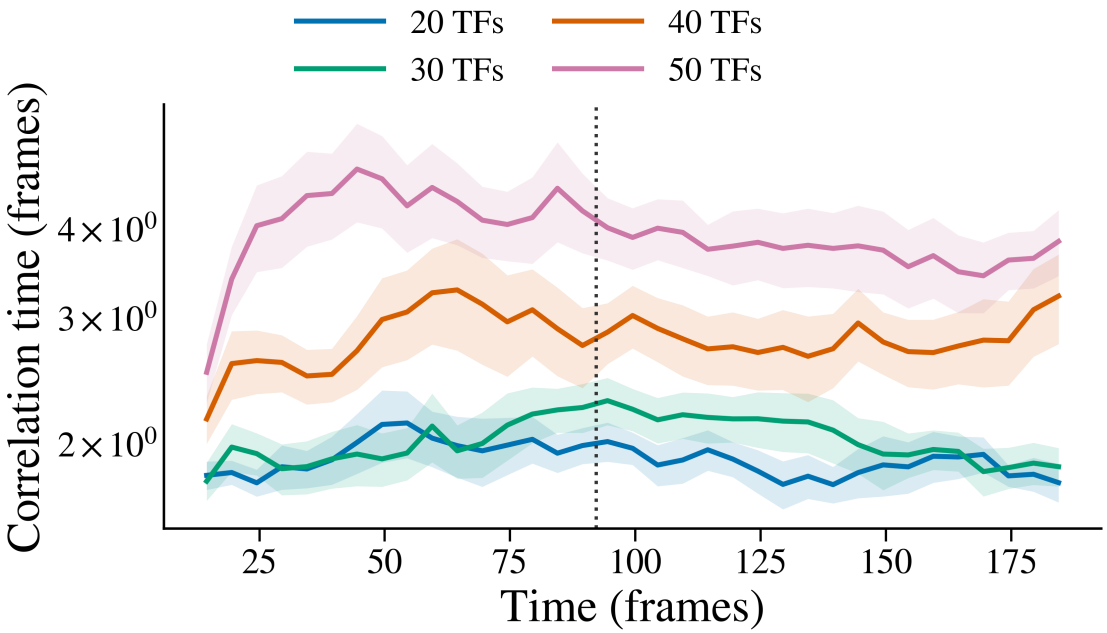

Figure 16. Correlation time τ_{corr} (log scale) over the full hysteresis trajectory for different TF copy numbers. The vertical dotted line marks the start of reversal. Higher TF number increases τ_{corr} and sustains long correlation times into the reverse phase, indicating stronger kinetic persistence.

Figure 16 shows $\tau_{\mathrm{corr}}(t)$ over the full hysteresis trajectory for different TF numbers. The separation in τ_{corr} is already visible at early times, before reversal begins. This indicates that TF abundance affects the baseline persistence of TU activity states in the stem-like regime, likely through more stable local chromatin–TF associations and more persistent clustering. During the forward differentiation phase, τ_{corr} increases in all cases, consistent with the suppression of OFF→ON and ON→OFF switching identified in Fig. 14(b). However, the magnitude of this increase depends strongly on TF number: higher TF abundance produces substantially longer correlation times, indicating that TU activity states become more persistent when TFs are abundant.

The crucial observation comes after reversal begins, marked by the vertical line in Fig. 16. At low TF number, τ_{corr} decreases towards its early-time values, indicating partial recovery of faster switching. At high TF number, by contrast, τ_{corr} remains elevated long after reversal begins. In other words, the system retains slow switching and long-lived ON/OFF states even when the differentiation protocol is undone. This is the first indication of canalisation: the differentiated state acquires transcriptional stability that is not immediately erased by reversing the external rules.

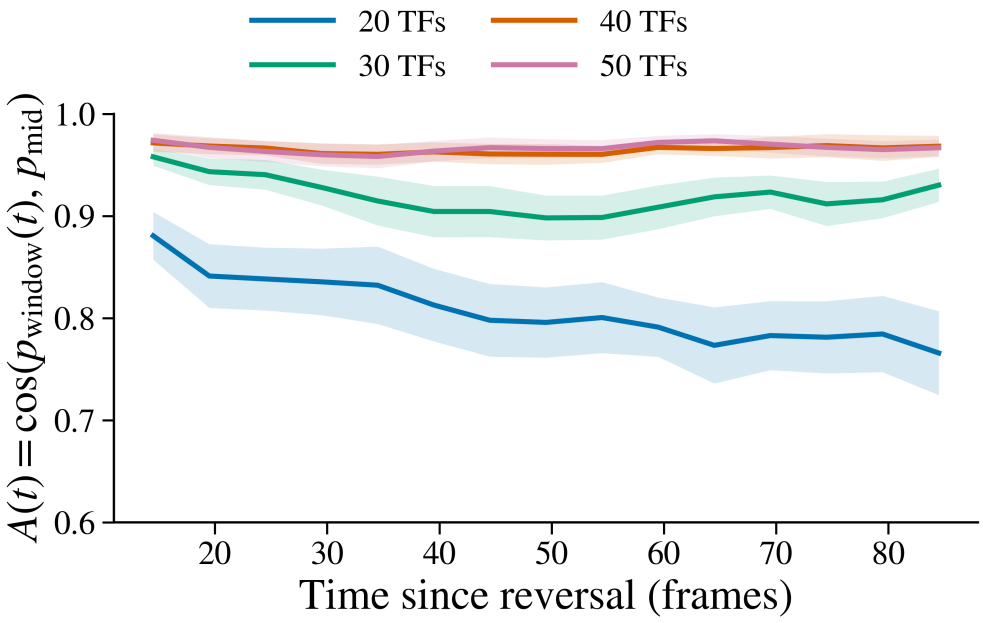

Figure 17. Differentiated-state similarity during the reversal phase. $A(t)$ is the cosine similarity between the window-averaged programme vector and the mid-trajectory differentiated reference. High TF number retains strong similarity (strong memory), whereas low TF number shows pronounced decay (reprogramming / weak memory).

While τ_{corr} quantifies the local persistence of TU activity states, it does not by itself tell us whether the *pattern* of activity returns towards the stem-like state. To address this, we analyse the transcriptional programme vector $\mathbf{p}(t)$ defined in Sec. 4.2 and compute its cosine similarity to the mid-trajectory differentiated reference $\mathbf{p}_{\mathrm{mid}}$. During the reversal phase, we therefore evaluate

$$A(t) = \cos\big(\mathbf{p}(t),\, \mathbf{p}_{\mathrm{mid}}\big),$$

where $\mathbf{p}(t)$ is the window-averaged activity vector. By construction, $A(t) \approx 1$ indicates strong retention of the differentiated transcriptional configuration, whereas smaller values indicate drift away from that state.

Figure 17 shows that TF number strongly controls this reversibility. At high TF abundance (40–50 TFs per type), $A(t)$ remains close to unity throughout the reversal phase. The transcriptional pattern therefore remains aligned with the differentiated reference even as the TF-collapse protocol is undone. In contrast, at low TF abundance (20 TFs per type), $A(t)$

decreases substantially after reversal begins. The system drifts away from the differentiated programme and explores configurations that are less aligned with it.

Together, Figs. 16 and 17 show that memory has two aspects: local persistence, reflected in slow switching, and global persistence, reflected in retention of the differentiated transcriptional programme. Higher TF number strengthens both.

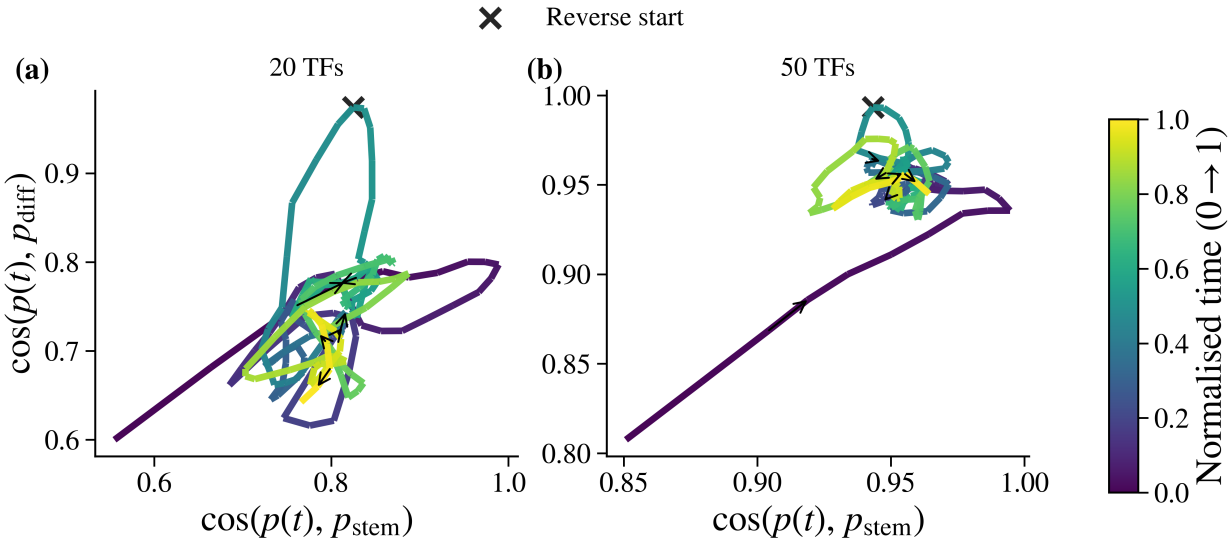

Figure 18. Phase-space trajectories in transcriptional identity space for 20 TFs (a) and 50 TFs (b). Axes show cosine similarity to stem and differentiated references; colour encodes normalised time, and the cross marks the reversal start. High TF number produces a compact differentiated attractor that persists through reversal, whereas low TF number exhibits broader exploration and greater departure from the differentiated region.

To visualise these effects geometrically, and to begin interpreting the differentiated state as an attractor, we represent each trajectory in the two-dimensional state space

$$x(t) = \cos\big(\mathbf{p}(t), \mathbf{p}_{\text{stem}}\big), \qquad y(t) = \cos\big(\mathbf{p}(t), \mathbf{p}_{\text{diff}}\big).$$

In this space, forward differentiation corresponds to motion away from the stem reference and towards the differentiated reference. Perfect reversibility would imply that the reverse trajectory retraces this path, although exact retracing is not expected in a stochastic system.

Figure 18 compares 20 TFs and 50 TFs. At high TF number, the trajectory rapidly approaches a compact region strongly aligned with the differentiated reference and remains confined there during reversal. There is little systematic movement back towards the stem-like direction. This confined, attractor-like behaviour is characteristic of strong canalisation: once established, the differentiated configuration behaves as a deep basin in state space. At low TF number, by contrast, the trajectory is broader and exhibits a clearer drift away from the differentiated region during reversal. The system does not remain tightly confined, and this broader exploration is consistent with weaker retention of the differentiated state and a shallower basin.

Taken together, the three measurements give a coherent picture. Correlation times (Fig. 16) show that high TF number produces long-lived TU activity states that persist even after reversal. Similarity to the differentiated programme (Fig. 17) shows that the overall transcriptional pattern is retained at high TF number but partially lost at low TF number. Finally, phase-space trajectories (Fig. 18) show that the differentiated state behaves like an attractor at high TF number, whereas at lower TF number the system departs from that state more readily.

In the Waddington picture, TF abundance therefore acts as a control knob that tunes the effective depth of the differentiated valley. Increasing TF number deepens the basin and strengthens memory, without requiring additional chromatin-accessibility constraints. Decreasing TF number leaves the basin shallower and more reversible. Thus, while TF number only weakly affects the qualitative features of fossilisation and commitment during differentiation, it plays a central role in determining how strongly cellular identity is remembered once established.

This attractor-like behaviour is consistent with several theoretical descriptions of cell-fate commitment. Ref. [72], for example, showed that differentiation can be understood as a transition from a multipotent attractor into deeper differentiated basins separated by asymmetric barriers. In that framework, barrier height controls transition times, so that forward transitions are favoured over reverse transitions, giving rise to an effective arrow of time. However, their model is formulated at the level of a minimal two-gene circuit and does not address transcriptional bursting, chromatin organisation, or dynamic hysteresis under reversal. Similarly, [73] analyse HL60 differentiation using an attractor-like description of gene-regulatory dynamics, identifying stable cell-fate states and characterising their relative robustness in state space, but without quantifying reversibility through an explicit hysteresis protocol or linking stability to correlation times or bursting statistics. Finally, [74] construct a deterministic mapping of Waddington's epigenetic landscape from gene-regulatory equations, explicitly associating differentiated identities with basin depth and stability, but again without incorporating stochastic switching times or transcriptional correlation measures.

Our results do not claim to reproduce these landscape constructions directly. Rather, they provide a complementary realisation of the same attractor picture using a polymer–TF bridging model. The conclusions drawn here should therefore be interpreted within the specific dynamical regime explored by this model, rather than as universal statements about all possible mechanisms of cellular memory or canalisation. By quantifying memory through telegraph correlation times, similarity of transcriptional programmes, and explicit forward–reverse trajectories in an abstract state space, we make basin depth operational through transcriptional dynamics. The observation that increasing TF abundance sustains long correlation times and confines trajectories near the differentiated reference is therefore consistent with landscape interpretations of canalisation [72, 73, 74], while extending them by linking attractor stability directly to transcriptional dynamics and TF-mediated chromatin organisation.

5 Conclusion and Future Prospects

In this project, we asked how chromatin organisation reshapes transcriptional dynamics across Waddington's landscape. Using a coarse-grained polymer model of bridging-induced phase separation, we represented chromatin as a connected bead–spring fibre containing transcription units (TUs), and transcription factors (TFs) as diffusing binders that switch stochastically between binding and non-binding states. From the resulting 3D chromatin–TF dynamics, we constructed binary TU activity traces and analysed them using a telegraph model to extract effective switching rates, burst statistics, and correlation times. This provided a minimal physical framework in which chromatin organisation, TF clustering, and transcriptional activity could be studied together, and in which the abstract language of Waddington's landscape could be related to explicit dynamical observables.

Our static analysis compared fixed chromatin states — stem-like, differentiated, and senescent — and established the phenomenon of fossilisation. In the stem-like state, TU mean activities span a broad intermediate range, consistent with flexible switching. In differentiated and senescent states, by contrast, intermediate regimes are depleted and activity accumulates near extreme ON or OFF values, producing a larger locked fraction. Rank–activity curves and cumulative-activity distributions further showed that transcription becomes concentrated into a smaller subset of dominant TUs, even when the global mean output changes only modestly. Burst frequencies shift downward at the same time, indicating that fossilisation is associated primarily with suppressed burst initiation. Differentiation and senescence therefore reshape transcription less by uniformly changing total activity than by redistributing activity away from intermediate switching regimes and towards effectively locked configurations. In the language of the landscape, these results are consistent with a progressive loss of accessible dynamical states.

We then moved from static endpoints to dynamical differentiation, allowing the system to descend Waddington's landscape in two distinct ways: chromatin remodelling through A/B polymer organisation, and TF colour collapse through a reduction in regulatory diversity. Both protocols drive progressive drift away from the initial stem-like programme and produce a strong increase in heterogeneity, confirming that activity becomes increasingly unevenly distributed across TUs. At the same time, the two routes separate different aspects of landscape descent. Chromatin remodelling yields a more stratified final state and a stronger late-time reduction in plasticity, whereas TF collapse produces a larger overall programme drift while retaining more residual plasticity. Together, these results support a two-stage picture of differentiation in the model: an early exploratory stage, in which different runs settle into different active subsets as the landscape is reshaped, followed by a later commitment stage, in which plasticity decreases and the intermediate TU population is depleted as more units become locked. Telegraph inference shows that this commitment is driven mainly by suppression of OFF$\rightarrow$ON transitions, while changes in ON$\rightarrow$OFF switching are weaker. Physically, this means that inactive states become progressively harder to escape, so differentiation acts chiefly by restricting re-entry into active configurations.

Finally, we quantified memory through a hysteresis protocol in which TF-collapse differentiation is reversed and the degree of retained transcriptional identity is measured. Here TF

abundance acts as a control parameter for canalisation. Higher TF number produces longer correlation times, preserves stronger similarity to the differentiated programme during reversal, and yields more compact phase-space trajectories that remain close to the differentiated reference rather than retracing towards a stem-like path. Lower TF number shows greater recovery of switching, a larger decay in similarity, and broader exploration during reversal. In the Waddington picture, TF abundance therefore tunes the effective depth of the differentiated basin: higher TF number deepens the basin and strengthens memory, whereas lower TF number leaves it shallower and more reversible. Within the model, this identifies a simple physical route by which TF abundance can stabilise cell-fate states without requiring a fully detailed biochemical description of epigenetic regulation.

These conclusions should, however, be interpreted within the specific dynamical regime studied here. The parameter values listed in Table 2 were chosen to place the system in a physically reasonable and sufficiently dynamic regime in which bridge-induced clustering, switching, and arrested coarsening could all be resolved on accessible simulation timescales. They were not varied systematically, and the aim of this work was not to establish universality across parameter space. The results therefore show what can arise robustly within this model and parameter regime, but they do not by themselves determine how strongly quantities such as memory, cluster competition, or transcriptional locking depend on the microscopic interaction strengths, switching rules, or sampling timescales.

This also points directly to the most important directions for future work. First, the robustness of fossilisation, commitment, and canalisation should be tested by varying the underlying simulation parameters more systematically, particularly the interaction strengths, TF abundance, switching rate, and chromatin-state protocols. Second, the structural basis of memory should be analysed more explicitly by connecting telegraph observables — TU locking, burst suppression, and correlation times — to measurable chromatin features such as local density, TF–TU bridge persistence, contact enrichment, cluster formation, and structural diversity [75]. Third, it would be valuable to examine how much additional information is present on timescales finer than the saved-frame resolution used here, for example by comparing frame-level telegraph statistics with the intermediate TF switching interval and the shorter-timescale rearrangement of local chromatin–TF contacts. This would help clarify which aspects of 3D organisation are merely correlated with transcriptional memory and which are most important for sustaining it.

More broadly, the programme vector defined in Eq. 24 may provide a route towards a more explicit dynamical mapping of Waddington's landscape, for example by constructing flow fields in programme space analogous to approaches developed for single-cell data [76, 77]. Associating transcriptional and structural observables — such as burst frequency, noise, contact structure, and cluster diversity — with positions in this space could help identify the dominant variables controlling developmental progression and thereby reduce the effective dimensionality of the problem. In this sense, the present work provides both a concrete result and a starting point: even a minimal physical model can already reproduce key features associated with cellular development, namely the loss of intermediate states, suppressed reactivation, and memory. This makes it a useful framework for future work aimed at connecting chromatin physics more directly to real biological systems.

References

[1] Bruce Alberts, Rebecca Heald, Alexander Johnson, David Morgan, Martin Raff, Keith Roberts, and Peter Walter. *Molecular Biology of the Cell*. W. W. Norton & Company, 7 edition, 2022.

[2] Scott F. Gilbert. *Developmental Biology*. Sinauer Associates, 10 edition, 2014.

[3] Sharmistha Kundu and Craig L Peterson. Role of chromatin states in transcriptional memory. *Biochimica et Biophysica Acta (BBA)-General Subjects*, 1790(6):445–455, 2009.

[4] Conrad Hal Waddington. *The Strategy of the Genes*. Routledge, 2014.

[5] Adrian Bird. Perceptions of epigenetics. *Nature*, 447(7143):396–398, 2007.

[6] Tom Misteli. Beyond the sequence: Cellular organization of genome function. *Cell*, 128(4):787–800, 2007.

[7] Job Dekker and Leonid Mirny. The 3D genome as moderator of chromosomal communication. *Cell*, 164(6):1110–1121, 2016.

[8] Boyan Bonev and Giacomo Cavalli. Organization and function of the 3D genome. *Nature Reviews Genetics*, 17(11):661–678, 2016.

[9] George Papadogkonas, Dionysios-Alexandros Papamatheakis, and Charalampos Spilianakis. 3D genome organization as an epigenetic determinant of transcription regulation in T cells. *Frontiers in Immunology*, 13:921375, 2022.

[10] Eran Meshorer and Tom Misteli. Chromatin in pluripotent embryonic stem cells and differentiation. *Nature Reviews Molecular Cell Biology*, 7(7):540–546, 2006.

[11] Jesse R. Dixon, Inkyung Jung, Siddarth Selvaraj, Yin Shen, Jessica E. Antosiewicz-Bourget, Ah Young Lee, Zhen Ye, Audrey Kim, Nisha Rajagopal, Wei Xie, et al. Chromatin architecture reorganization during stem cell differentiation. *Nature*, 518(7539):331–336, 2015.

[12] Tamir Chandra, Philip Andrew Ewels, Stefan Schoenfelder, Mayra Furlan-Magaril, Steven William Wingett, Kristina Kirschner, Jean-Yves Thuret, Simon Andrews, Peter Fraser, and Wolf Reik. Global reorganization of the nuclear landscape in senescent cells. *Cell Reports*, 10(4):471–483, 2015.

[13] Haitham A. Shaban and Susan M. Gasser. Dynamic 3D genome reorganization during senescence: Defining cell states through chromatin. *Cell Death & Differentiation*, 32(1):9–15, 2025.

[14] Chris A. Brackley, Stephen Taylor, Argyris Papantonis, Peter R. Cook, and Davide Marenduzzo. Nonspecific bridging-induced attraction drives clustering of DNA-binding proteins and genome organization. *Proceedings of the National Academy of Sciences*, 110(38):E3605–E3611, 2013.

[15] Chris A. Brackley, Benno Liebchen, Davide Michieletto, Francois Mouvet, Peter R. Cook, and Davide Marenduzzo. Ephemeral protein binding to DNA shapes stable nuclear bodies and chromatin domains. *Biophysical Journal*, 112(6):1085–1093, 2017.

[16] Chris A. Brackley. Polymer compaction and bridging-induced clustering of protein-inspired patchy particles. *Journal of Physics: Condensed Matter*, 32(31):314002, 2020.

[17] Shuvadip Dutta and Ranjith Padinhateeri. Polymer simulations of chromatin: Connecting 3D organization and dynamics to function. *Current Opinion in Genetics & Development*, 96:102429, 2026.

[18] Anders S. Hansen, Claudia Cattoglio, Xavier Darzacq, and Robert Tjian. Recent evidence that TADs and chromatin loops are dynamic structures. *Nucleus*, 9(1):20–32, 2018.

[19] Job Dekker and Leonid A. Mirny. The chromosome folding problem and how cells solve it. *Cell*, 187(23):6424–6450, 2024.

[20] Jeroen S. van Zon, Marco J. Morelli, Sorin Tănase-Nicola, and Pieter Rein ten Wolde. Diffusion of transcription factors can drastically enhance the noise in gene expression. *Biophysical Journal*, 91(12):4350–4367, 2006.

[21] Rama Reddy Goluguri, Mourad Sadqi, and Victor Munoz. Transcriptional factors control their diffusion on DNA by modulating their dynamics. *Biophysical Journal*, 121(3):457a, 2022.

[22] Zhixing Cao, Tatiana Filatova, Diego A. Oyarzún, and Ramon Grima. A stochastic model of gene expression with polymerase recruitment and pause release. *Biophysical Journal*, 119(5):1002–1014, 2020.

[23] Songhao Luo, Zhenquan Zhang, Zihao Wang, Xiyan Yang, Xiaoxuan Chen, Tianshou Zhou, and Jiajun Zhang. Inferring transcriptional bursting kinetics from single-cell snapshot data using a generalized telegraph model. *Royal Society Open Science*, 10(4):221057, 2023.

[24] Feng Jiao, Jing Li, Ting Liu, Yifeng Zhu, Wenhao Che, Leonidas Bleris, and Chen Jia. What can we learn when fitting a simple telegraph model to a complex gene expression model? *PLOS Computational Biology*, 20(5):e1012118, 2024.

[25] Sui Huang. Reprogramming cell fates: Reconciling rarity with robustness. *BioEssays*, 31(5):546–560, 2009.

[26] Katherine M. Aird and Rugang Zhang. Detection of senescence-associated heterochromatin foci (SAHF). In *Cell Senescence: Methods and Protocols*, pages 185–196. Springer, 2012.

[27] Masashi Narita, Sabrina Nuñez, Edith Heard, Masako Narita, Athena W. Lin, Stephen A. Hearn, David L. Spector, Gregory J. Hannon, and Scott W. Lowe. Rb-mediated heterochromatin formation and silencing of E2F target genes during cellular senescence. *Cell*, 113(6):703–716, 2003.

[28] Masashi Narita, Masako Narita, Valery Krizhanovsky, Sabrina Nuñez, Agustin Chicas, Stephen A. Hearn, Michael P. Myers, and Scott W. Lowe. A novel role for high-mobility group a proteins in cellular senescence and heterochromatin formation. *Cell*, 126(3):503–514, 2006.

[29] Anthony A. Hyman, Christoph A. Weber, and Frank Jülicher. Liquid-liquid phase separation in biology. *Annual Review of Cell and Developmental Biology*, 30(1):39–58, 2014.

[30] Amy R. Strom, Alexander V. Emelyanov, Mustafa Mir, Dmitry V. Fyodorov, Xavier Darzacq, and Gary H. Karpen. Phase separation drives heterochromatin domain formation. *Nature*, 547(7662):241–245, 2017.

[31] Adam G. Larson, Daniel Elnatan, Madeline M. Keenen, Michael J. Trnka, Jonathan B. Johnston, Alma L. Burlingame, David A. Agard, Sy Redding, and Geeta J. Narlikar. Liquid droplet formation by HP1α suggests a role for phase separation in heterochromatin. *Nature*, 547(7662):236–240, 2017.

[32] Leonid A. Mirny, Maxim Imakaev, and Nezar Abdennur. Two major mechanisms of chromosome organization. *Current Opinion in Cell Biology*, 58:142–152, 2019.

[33] Fabian Erdel and Karsten Rippe. Formation of chromatin subcompartments by phase separation. *Biophysical Journal*, 114(10):2262–2270, 2018.

[34] Mark D. Biggin. Animal transcription networks as highly connected, quantitative continua. *Developmental Cell*, 21(4):611–626, 2011.

[35] Lyubomira Chakalova and Peter Fraser. Organization of transcription. *Cold Spring Harbor Perspectives in Biology*, 2(9):a000729, 2010.

[36] Argyris Papantonis and Peter R. Cook. Transcription factories: Genome organization and gene regulation. *Chemical Reviews*, 113(11):8683–8705, 2013.

[37] Mariano Barbieri, Mita Chotalia, James Fraser, Liron-Mark Lavitas, Josée Dostie, Ana Pombo, and Mario Nicodemi. Complexity of chromatin folding is captured by the strings and binders switch model. *Proceedings of the National Academy of Sciences*, 109(40):16173–16178, 2012.

[38] Renko de Vries. Influence of mobile DNA-protein-DNA bridges on DNA configurations: Coarse-grained monte carlo simulations. *The Journal of Chemical Physics*, 135(12), 2011.

[39] Miriam Fritsche, Songling Li, Dieter W. Heermann, and Paul A. Wiggins. A model for Escherichia coli chromosome packaging supports transcription factor-induced DNA domain formation. *Nucleic Acids Research*, 40(3):972–980, 2012.

[40] Reza Kalhor, Harianto Tjong, Nimanthi Jayathilaka, Frank Alber, and Lin Chen. Genome architectures revealed by tethered chromosome conformation capture and population-based modeling. *Nature Biotechnology*, 30(1):90–98, 2012.

[41] Davide Marenduzzo, Cristian Micheletti, and Peter R. Cook. Entropy-driven genome organization. *Biophysical Journal*, 90(10):3712–3721, 2006.

[42] Mariliis Tark-Dame, Martijn S. Luijsterburg, Dieter W. Heermann, and Roel van Driel. Of chromatin folding and chromatin-associated processes. *Genome Organization and Function in the Cell Nucleus*, page 535, 2011.

[43] Paul A. Wiggins, Remus T. Dame, Maarten C. Noom, and Gijs J. L. Wuite. Protein-mediated molecular bridging: A key mechanism in biopolymer organization. *Biophysical Journal*, 97(7):1997–2003, 2009.

[44] Remus T. Dame, Maarten C. Noom, and Gijs J. L. Wuite. Bacterial chromatin organization by H-NS protein unravelled using dual DNA manipulation. *Nature*, 444(7117):387–390, 2006.

[45] Anatoly A. Zinchenko, Kenichi Yoshikawa, and Damien Baigl. Compaction of single-chain DNA by histone-inspired nanoparticles. *Physical Review Letters*, 95(22):228101, 2005.

[46] Anatoly A. Zinchenko, Takahiro Sakaue, Sumiko Araki, Kenichi Yoshikawa, and Damien Baigl. Single-chain compaction of long duplex DNA by cationic nanoparticles: Modes of interaction and comparison with chromatin. *The Journal of Physical Chemistry B*, 111(11):3019–3031, 2007.

[47] Je-Kyung Ryu, Céline Bouchoux, Hon Wing Liu, Eugene Kim, Masashi Minamino, Ralph de Groot, Allard J. Katan, Andrea Bonato, Davide Marenduzzo, Davide Michieletto, et al. Bridging-induced phase separation induced by cohesin SMC protein complexes. *Science Advances*, 7(7):eabe5905, 2021.

[48] Sinan Kilic, Andreas L. Bachmann, Louise C. Bryan, and Beat Fierz. Multivalency governs HP1α association dynamics with the silent chromatin state. *Nature Communications*, 6(1):7313, 2015.

[49] Peter R. Cook and Davide Marenduzzo. Transcription-driven genome organization: A model for chromosome structure and the regulation of gene expression tested through simulations. *Nucleic Acids Research*, 46(19):9895–9906, 2018.

[50] Patrick Cramer. Organization and regulation of gene transcription. *Nature*, 573(7772):45–54, 2019.

[51] Peter R. Cook. The organization of replication and transcription. *Science*, 284(5421):1790–1795, 1999.

[52] Argyris Papantonis, Takahide Kohro, Sabyasachi Baboo, Joshua D. Larkin, Binwei Deng, Patrick Short, Shuichi Tsutsumi, Stephen Taylor, Yasuharu Kanki, Mika Kobayashi, et al. TNFα signals through specialized factories where responsive coding and mirna genes are transcribed. *The EMBO Journal*, 31(23):4404, 2012.

[53] Krishna Shrinivas, Benjamin R. Sabari, Eliot L. Coffey, Isaac A. Klein, Ann Boija, Alicia V. Zamudio, Jurian Schuijers, Nancy M. Hannett, Phillip A. Sharp, Richard A. Young, et al. Enhancer features that drive formation of transcriptional condensates. *Molecular Cell*, 75(3):549–561, 2019.

[54] Chris A. Brackley, Nick Gilbert, Davide Michieletto, Argyris Papantonis, M. C. F. Pereira, P. R. Cook, and Davide Marenduzzo. Complex small-world regulatory networks emerge from the 3D organisation of the human genome. *Nature Communications*, 12(1):5756, 2021.

[55] Alexander J Stewart, Sridhar Hannenhalli, and Joshua B Plotkin. Why transcription factor binding sites are ten nucleotides long. *Genetics*, 192(3):973–985, 2012.

[56] Andrew J. Bannister and Tony Kouzarides. Regulation of chromatin by histone modifications. *Cell Research*, 21(3):381–395, 2011.

[57] Howard Cedar and Yehudit Bergman. Programming of DNA methylation patterns. *Annual Review of Biochemistry*, 81(1):97–117, 2012.

[58] Yi-Ping Wang and Qun-Ying Lei. Metabolic recoding of epigenetics in cancer. *Cancer Communications*, 38(1):25, 2018.

[59] Julie Borgel, Sylvain Guibert, Yufeng Li, Hatsune Chiba, Dirk Schübeler, Hiroyuki Sasaki, Thierry Forné, and Michael Weber. Targets and dynamics of promoter DNA methylation during early mouse development. *Nature Genetics*, 42(12):1093–1100, 2010.

[60] Sunny Sun-Kin Chan and Michael Kyba. What is a master regulator? *Journal of Stem Cell Research & Therapy*, 3:114, 2013.

[61] Stephen J. Tapscott, Robert L. Davis, Mathew J. Thayer, Pei-Feng Cheng, Harold Weintraub, and Andrew B. Lassar. MyoD1: A nuclear phosphoprotein requiring a myc homology region to convert fibroblasts to myoblasts. *Science*, 242(4877):405–411, 1988.

[62] Lorraine Robb, Ngaire J. Elwood, Andrew G. Elefanty, Frank Köntgen, Ruili Li, Louise D. Barnett, and C. Glenn Begley. The scl gene product is required for the generation of all hematopoietic lineages in the adult mouse. *The EMBO Journal*, 15(16):4123–4129, 1996.

[63] Catherine Porcher, Wojciech Swat, Karen Rockwell, Yuko Fujiwara, Frederick W. Alt, and Stuart H. Orkin. The T-cell leukemia oncoprotein SCL/tal-1 is essential for development of all hematopoietic lineages. *Cell*, 86(1):47–57, 1996.

[64] Pu Zhang, Gerhard Behre, Jing Pan, Atsushi Iwama, Nawarat Wara-Aswapati, Hanna S. Radomska, Philip E. Auron, Daniel G. Tenen, and Zijie Sun. Negative cross-talk between hematopoietic regulators: GATA proteins repress PU.1. *Proceedings of the National Academy of Sciences*, 96(15):8705–8710, 1999.

[65] Pu Zhang, Xiaobo Zhang, Atsushi Iwama, Channing Yu, Kent A. Smith, Beatrice U. Mueller, Salaija Narravula, Bruce E. Torbett, Stuart H. Orkin, and Daniel G. Tenen. PU.1 inhibits GATA-1 function and erythroid differentiation by blocking GATA-1 DNA binding. *Blood*, 96(8):2641–2648, 2000.

[66] Rugang Zhang, Maxim V. Poustovoitov, Xiaofen Ye, Hidelita A. Santos, Wei Chen, Sally M. Daganzo, Jan P. Erzberger, Ilya G. Serebriiskii, Adrian A. Canutescu, Roland L. Dunbrack, et al. Formation of MacroH2A-containing senescence-associated heterochromatin foci and senescence driven by ASF1a and HIRA. *Developmental Cell*, 8(1):19–30, 2005.

[67] Stefan Semrau, Johanna E. Goldmann, Magali Soumillon, Tarjei S. Mikkelsen, Rudolf Jaenisch, and Alexander Van Oudenaarden. Dynamics of lineage commitment revealed by single-cell transcriptomics of differentiating embryonic stem cells. *Nature Communications*, 8(1):1096, 2017.

[68] Ian R. McCracken, Richard S. Taylor, Fatma O. Kok, Fernando de la Cuesta, Ross Dobie, Beth E. P. Henderson, Joanne C. Mountford, Axelle Caudrillier, Neil C. Henderson, Chris P. Ponting, et al. Transcriptional dynamics of pluripotent stem cell-derived endothelial cell differentiation revealed by single-cell RNA sequencing. *European Heart Journal*, 41(9):1024–1036, 2020.

[69] Reem Elorbany, Joshua M. Popp, Katherine Rhodes, Benjamin J. Strober, Kenneth Barr, Guanghao Qi, Yoav Gilad, and Alexis Battle. Single-cell sequencing reveals lineage-specific dynamic genetic regulation of gene expression during human cardiomyocyte differentiation. *PLOS Genetics*, 18(1):e1009666, 2022.

[70] Caleb Weinreb, Alejo Rodriguez-Fraticelli, Fernando D. Camargo, and Allon M. Klein. Lineage tracing on transcriptional landscapes links state to fate during differentiation. *Science*, 367(6479):eaaw3381, 2020.

[71] Kazutoshi Takahashi and Shinya Yamanaka. Induction of pluripotent stem cells from mouse embryonic and adult fibroblast cultures by defined factors. *Cell*, 126(4):663–676, 2006.

[72] Jin Wang, Li Xu, Erkang Wang, and Sui Huang. The potential landscape of genetic circuits imposes the arrow of time in stem cell differentiation. *Biophysical Journal*, 99(1):29–39, 2010.

[73] Albert C. Huang, Limei Hu, Stuart A. Kauffman, Wei Zhang, and Ilya Shmulevich. Using cell fate attractors to uncover transcriptional regulation of HL60 neutrophil differentiation. *BMC Systems Biology*, 3(1):20, 2009.

[74] Sudin Bhattacharya, Qiang Zhang, and Melvin E. Andersen. A deterministic map of waddington's epigenetic landscape for cell fate specification. *BMC Systems Biology*, 5(1):85, 2011.

[75] Michael Chiang, Chris A. Brackley, Catherine Naughton, Ryu-Suke Nozawa, Cleis Battaglia, Davide Marenduzzo, and Nick Gilbert. Genome-wide chromosome architecture prediction reveals biophysical principles underlying gene structure. *Cell Genomics*, 4(12):100698, 2024.

[76] Ligang Zhu, Songlin Yang, Kun Zhang, Hong Wang, Xiaona Fang, and Jin Wang. Uncovering underlying physical principles and driving forces of cell differentiation and reprogramming from single-cell transcriptomics. *Proceedings of the National Academy of Sciences*, 121(34):e2401540121, 2024.

[77] Gioele La Manno, Ruslan Soldatov, Amit Zeisel, Emelie Braun, Hannah Hochgerner, Viktor Petukhov, Katja Lidschreiber, Maria E. Kastriti, Peter Lönnerberg, Alessandro Furlan, et al. RNA velocity of single cells. *Nature*, 560(7719):494–498, 2018.