



Latent space of a small genetic network: Geometry of dynamics and information

Rabea Seyboldt^a, Juliette Lavoie^a, Adrien Henry^a, Jules Vanaret^{a,b}, Mariela D. Petkova^c, Thomas Gregor^{d,e,f}, and Paul François^{a,1}

Edited by Eric Siggia, The Rockefeller University, New York, NY; received July 23, 2021; accepted March 31, 2022

The high-dimensional character of most biological systems presents genuine challenges for modeling and prediction. Here we propose a neural network–based approach for dimensionality reduction and analysis of biological gene expression data, using, as a case study, a well-known genetic network in the early *Drosophila* embryo, the gap gene patterning system. We build an autoencoder compressing the dynamics of spatial gap gene expression into a two-dimensional (2D) latent map. The resulting 2D dynamics suggests an almost linear model, with a small bare set of essential interactions. Maternally defined spatial modes control gap genes positioning, without the classically assumed intricate set of repressive gap gene interactions. This, surprisingly, predicts minimal changes of neighboring gap domains when knocking out gap genes, consistent with previous observations. Latent space geometries in maternal mutants are also consistent with the existence of such spatial modes. Finally, we show how positional information is well defined and interpretable as a polar angle in latent space. Our work illustrates how optimization of small neural networks on medium-sized biological datasets is sufficiently informative to capture essential underlying mechanisms of network function.

systems biology | developmental biology | machine learning | dimensionality reduction | *Drosophila* gap genes

The dimensionality of models in biology often mirrors the vast array of molecular interactions underlying life processes (1), making it difficult to extract the most significant information about the system and general principles (2, 3). Recent advances in machine learning (4) allow for a new class of emerging strategies to reduce a system to its bare modes, often in the form of self-generating neural networks that mimic the dynamics of the original system (see, e.g., ref. 5 for physical systems). The attraction for low-dimensional models is their often straightforward interpretability, making it easier to identify general principles that are common to entire classes of systems (6–12), which can be put to experimental verification. The challenge in applying this approach to biological systems is the comparatively small number of data, which might rapidly lead to overfitting (13), and it is generally unclear how to systematically derive small models. Here we present a case study where we reduce the complex spatiotemporal dynamics of a well-studied gene regulatory network into a low-dimensional representation, to derive a general understanding of the core features of the network dynamics and its function.

During the initial stages of fruit fly development, the early embryo presents an experimentally accessible system for learning about genetic networks (14–18). In particular, the both experimentally and theoretically well-studied gap gene network, involved in segment patterning along the anterior–posterior (AP) axis of the embryo, provides an ideal example for the dimensionality reduction problem. Gap genes form an interconnected layer in an otherwise feed-forward network that takes upstream inputs from primary maternal morphogens (Fig. 1A). Following early theoretical predictions by Meinhardt (19), it is generally proposed that gap genes are mutually repressive (17, 20–24). This defines an “alternating cushion” system, positioning gap genes’ expressions complementary to one another (25–27). Under downstream control of the gap gene layer are the pair-rule genes (28). Their expression occurs in stripes that are precisely and reproducibly positioned within the embryo, forming an outline for the segmented body plan of the fully developed organism (29). In this system, states of the network are the expression levels of the gap genes, and the functional information being encoded is the position of cells along the AP axis. This three-layered genetic cascade offers a textbook example of positional information (30), where morphogen concentrations eventually dictate local cellular fates. Biophysical approaches have been used to properly define, quantify, and reconstruct such information (31), unequivocally inferring cellular positions—and thus a proxy for cell fates—solely from gap gene concentrations (32).

The availability of high-precision, dynamical data (32), combined with previous theoretical studies (27, 33–35), led us to this case study for a data-driven, low-dimensionality neural network search using a small autoencoder (4, 36). Autoencoders are neural

Significance

In living systems all processes are intrinsically dynamic. But even the most basic biological dynamics are of such high-dimensional character that it is often difficult to deduce representations containing the most essential features with high predictive power. Here we consider the dynamics of a small genetic network driving early fly development and derive a picture representation that intuitively encodes the biological notion of positional information. We show how gene regulation and network circuitry are associated with geometric features in this picture, and how a parsimonious model for fly development separating time and space emerges naturally. Our work illustrates how small, informative representations of biological data serve for intuitive interpretation of complex biological regulation and dynamics.

Author contributions: R.S., A.H., and P.F. designed research; R.S., J.L., A.H., J.V., M.D.P., and P.F. performed research; R.S., J.L., A.H., J.V., and T.G. contributed new reagents/analytic tools; R.S., J.L., A.H., J.V., M.D.P., T.G., and P.F. analyzed data; and T.G. and P.F. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2022 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: paul.francois2@mcgill.ca.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2113651119/-DCSupplemental>.

Published June 22, 2022.




Fig. 1. Dimensionality reduction of gap gene dynamics using an autoencoder. (A) The gap gene system is part of the early *Drosophila* segmentation gene network, a cascade of patterning genes that specify cells along the embryo's AP axis. (Left) Cartoon expression profiles of maternal genes *bicoid*, *nanos*, and *torso*. (Middle) Gap gene expression for genes *giant*, *knirps*, *Krüppel*, and *hunchback*. (Right) Pair-rule genes expression. (B) Example of a trained nonlinear autoencoder with four input nodes (Left) for the four gap genes and compression into a layer with two nodes (Middle) using ReLU activation functions towards four output nodes (Right). Vertex colors indicate node weights after training (green for positive, red for negative weights). Node colors indicate the state of the network for an example with high *Kr* and *hb* input. (See *SI Appendix, section 1* for more detail.) (C) Encoding and decoding of the linear autoencoder. Encoding compresses four gap genes in two dimensions, realized by taking the differences of gene pairs. Decoding projects the two latent space dimensions onto the four gap genes. (D) Snapshot of Gaussian-filtered gap gene concentration profiles along the AP axis at $t = 30$ min into n.c. 14. Points of interest are labeled (1 to 4 for maxima, A/B for additional points). (E) Latent space profiles H1 and H2 for *D* with corresponding labels.

networks with a bottleneck layer containing only a few nodes. They are easy to train, as the output layer is directly fitted onto the input layer (36). If training is successful, the bottleneck layer encodes (almost) all the information contained in the data into a small number of nodes (37). The associated phase space defined by the bottleneck layer is generally called “latent space” (4). Contrary to more-complex dimensionality reduction techniques [such as t -distributed Stochastic Neighbor Embedding (t -SNE) (38) or Uniform Manifold Approximation and Projection (UMAP) (39)], the relationship between the full system and its low-dimensional latent space projection is mathematically explicit and intuitive. As we will show below, this allows us to go back and forth between the full and the latent space, thus offering a generalized natural route for low-dimension modeling of biological systems.

Applying this approach to the early fly embryo, we build an easy-to-interpret two-dimensional (2D) map to describe gap gene network expression levels in space and time. We identified simple spatiotemporal modes in latent space, pointing to a 2D model where time and space are separable: Gap gene positions are set up

by independent spatial modes, with a single common temporal mode. Mutual interactions only feature as weak perturbations. Data from maternal gene mutants are consistent with this view, revealing a modular nature of the wild-type (WT) latent space. Finally, the latent space representation provides an intuitive geometric interpretation of positional information. This analysis illustrates that small latent space models are helpful theoretical tools, which can be deterministically built from data to provide quantitative insights into high-dimensional biological systems.

Results

A Small Interpretable Autoencoder for Gap Gene Reconstruction. We use data from Petkova et al. (32) that provide expression levels of the gap genes *hunchback* (*hb*), *giant* (*gt*), *knirps* (*kni*), and *Krüppel* (*Kr*). They were measured simultaneously in nuclei located in the midsagittal plane in $N = 130$ WT embryos at different time points. Overall expression is projected along the assumed 1D major body axis (Fig. 1A) for each time point. Thus,

a data point corresponds to a 4D vector of expression levels (one dimension per gene), obtained for each position projected along the AP axis and for 60 time points corresponding to the 1-h duration of nuclear cycle (n.c.) 14.

To train autoencoders on these data, we tested multiple architectures with different numbers of intermediate layers, nodes, and nonlinearities. We typically used standard rectified linear units (ReLU) as activation functions ($\text{ReLU}(x) = x$ if $x > 0$, zero otherwise) (4). Autoencoders with a single node in the intermediate layer were insufficient to give adequate data descriptions, indicating that the spatiotemporal gap gene manifold is, at the very least, bidimensional. This is consistent with the assumption that at least two degrees of freedom are necessary, one for time and one for space. Conversely, when two nodes (H_1, H_2) are used in the intermediate layer, autoencoders are efficient at capturing/reconstructing most features of the gap gene dynamics (Fig. 1B). An example autoencoder reconstructs most features of the gap gene profiles, at all time points, such as peak positions and relative magnitude of gap genes (Fig. 1C and D). It misbehaves only by cutting off low gap gene concentrations, such as the late *kni* peak in the anterior (compare input and output in Fig. 1C). Increasing the number of intermediate layers did not lead to significant qualitative improvements, and the learned autoencoder structure is reproducible across training sets. Notice that the autoencoder does not “predict” any new data: It simply compresses the original data into the intermediate layer with two nodes, before reconstructing it.

After training, we observe a remarkable and transparent compartmentalization of the system.

- 1) While H_1 is solely a function of Gt and Kr, H_2 only depends on Kni and Hb. In both cases, gap genes contributed with opposing signs.
- 2) Likewise, in the output layer, Gt and Kr can be derived from H_1 only, and Kni and Hb from H_2 only, again with opposing signs.

Thus the gap gene system effectively reduces to a two-variable system that defines a 2D latent space. Such compression is possible because both variables, H_1 and H_2 , are controlled by mutually exclusive genes (e.g., positive H_1 directly encodes Gt, while negative H_1 encodes Kr). Our autoencoder recapitulates the “alternating cushions” (17, 27), which have been suggested to arise from mutual repression between nonoverlapping gap genes (i.e., respectively, (*hb, kni*) and (*gt, Kr*)).

Based on these results, we chose the simplest two-node autoencoder that handles all gap genes equally and significantly simplifies our analyses going forward, by defining

$$H_1(x, t) = \text{Kr}(x, t) - \text{Gt}(x, t), \quad [1]$$

$$H_2(x, t) = \text{Hb}(x, t) - \text{Kni}(x, t), \quad [2]$$

and the original data are reconstructed with the help of four ReLU functions,

$$\text{Gt} = \text{ReLU}(-H_1) \quad \text{Kr} = \text{ReLU}(H_1), \quad [3]$$

$$\text{Kni} = \text{ReLU}(-H_2) \quad \text{Hb} = \text{ReLU}(H_2). \quad [4]$$

Gap gene reconstruction by such an autoencoder is in excellent visual agreement between input data and decoded output (Fig. 1C and SI Appendix, section 1), perfectly capturing gap gene peaks and boundaries. Correspondence between averaged values of gap gene inputs and (H_1, H_2) is shown in Fig. 1D and E with

landmarks at different positions. We quantified and plotted correlations between data and reconstruction (SI Appendix, Fig. S4) and found an average error of less than 5% of the maximum gap gene values.

Analyzing the Dynamics of a Gene Network in Latent Space.

The dynamics of the 4D gap gene system, evolving in time and space, is thus compressed into a plane defined by two latent variables, H_1 and H_2 , with minimal loss of information. To gain insights about the dynamics and possible interactions among the gap genes, we consider $H_1(x, t)$ and $H_2(x, t)$ in latent space (Fig. 2A). At any given time, they map out a parametric curve (subsequently called “position manifold”) that represents positions x in latent space. The curve folds on itself in the upper quadrant of latent space (purple shading) and subsequently traces a roughly square box around the origin (Fig. 2B).

For the temporal progression of this position manifold, we identify two phases during the 1-h-long n.c. 14: the first before 45 min and the second after 45 min, corresponding to, respectively, a collective overall rise and fall of the profiles (Fig. 2C and D and Movie S1). Note that around $t = 45$ min is the time when positional information encoded by the gap genes is maximal (24), meaning that the dynamics leading up to this time encodes most of the functional information. In fact, during that first phase, the position manifold has a particularly simple shape: It mostly expands, and flow lines never cross (Fig. 2C and SI Appendix, Fig. S2), allowing for a considerable dimensional compression. By expressing the flow with differential equations

$$\frac{d}{dt}(H_1, H_2) = \mathcal{F}(H_1, H_2),$$

we effectively replace a 4×4 Gene Regulatory Network (GRN) by a 2×2 GRN in latent space. Since, in each part of the plane, each H_i encodes for a single gap gene, from a mathematical standpoint, local interactions between spatially adjacent gap genes, that is, (*gt, hb*), (*hb, Kr*), (*Kr, kni*), and (*kni, gt*), respectively, are sufficient to explain the dynamics of the position manifold up to $t = 45$ min.

The flow in latent space for $t < 45$ min (Fig. 2C) can be further decomposed into a dominant radial component and a weaker angular contribution. The latter corresponds to a very slow linear motion of gap gene expression peaks (i.e., at most, half a cell diameter per minute [SI Appendix, section 3 and Fig. S7]), directed toward the posterior boundary of the anterior *hb* domain. This motion is well known in the posterior (33, 40), but we notice that anterior gap genes move slightly toward the posterior (SI Appendix, Fig. S7), suggestive of a global systematic effect. We can mathematically remove this contribution, and the flow becomes radially symmetric in most of latent space (Fig. 2E and F). Importantly, if there are repressions between spatially adjacent gap genes so that one gene decreases the other one, one would observe angular components of the flow [e.g., on slow (41) or canalized manifolds (42)]. So a radial flow suggests an absence of repressions. We note that there is a strong remaining local angular component at the border of the (*gt, hb*) quadrant (arrow in Fig. 2E, landmark B in Fig. 1D) that could thus originate from a repressive interaction.

Flow Dynamics Implies Space–Time Separation. To understand the biophysical origin of the mostly radial flow, we resort to a standard Partial Differential Equations (PDE)-based toy model for four genes, produced in localized source regions along the AP axis (Fig. 3A and SI Appendix, section 6). For two of the simplest models, one with and the other one without mutual repression between

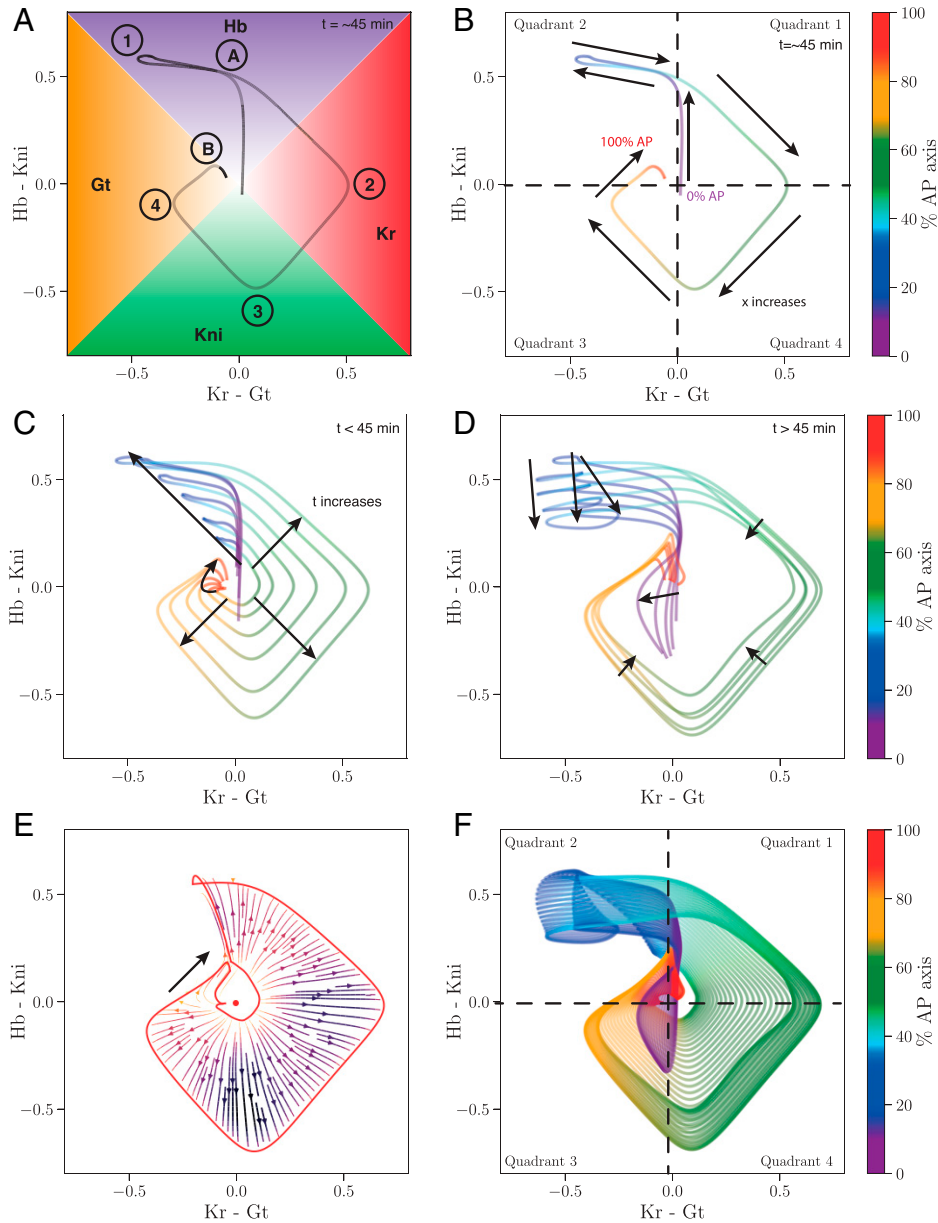


Fig. 2. Gap gene dynamics in latent space. (A) Snapshot of concentration along the AP axis in latent space (line) for $t = 45$ min. Points of interest are labeled as in Fig. 1D. Background color corresponds to the line color of the dominant gene (i.e., the gap gene with higher concentrations than all other ones) in each area. Throughout the text, we use the same latent space, where vertical direction corresponds to the *kni-hb* axis, and horizontal direction corresponds to the *gt-kr* axis. (B) Latent profile as in A, color coded by position along the AP axis. Arrows denote progression along increasing x values along the axis. (C) Latent profiles along the AP axis as in D, but for several time points with $t < 45$ min. For increasing time, the radius of the profiles increases (arrows). Note that, in quadrant 2, for $x > 80\%$ AP, a flow orthogonal to radial direction is observed. (D) Profiles for later time points with $t > 45$ min. In quadrants 1, 3, and 4, the radius decreases slightly (arrows); in quadrant 2, more complex dynamics with the creation of the fold structure is observed. (E) Flow profiles in latent space after speed removal for $t < 40$ min for positions $\leq 40\%$ AP. Borders of the flow and origin of latent space are shown in red. Notice that the flow is radial almost everywhere, except in quadrant 2, where there is an angular component parallel to the arrow. (F) Average of all gap gene data in 2D latent space. Each line corresponds to the position manifold at a different time point $t = [3, \dots, 53]$ min after start of cell cycle 14. Color encodes the position along the axis.

adjacent genes, the flow of the dynamics of the position manifolds in latent space is indeed radial (*SI Appendix, Figs. S13–S15*), initially forming concave lines with sharply pronounced corners, corresponding to the concentration maxima in the source regions. While the position manifold in latent space for the model without interactions eventually results in a square-shaped box (Fig. 3B), the model with interactions stabilizes in a concave shape (Fig. 3C). Such concavity is not observed in the gap gene system (Fig. 2E), which further excludes mutual repressions.

To explain the observed presence of a homogeneously expanding squared position manifold with a nonconcave shape at all times (even early), the simplest solution is to consider an

“adiabatic” version of this model, in which the system is effectively at steady state at all times (e.g., all production rates increasing slowly and uniformly under the influence of a global activator). Such a model can be written as

$$H_i(x, t) = \lambda(t)h_i(x), \quad [5]$$

defining separable spatial modes ($h_1(x)$ and $h_2(x)$) as well as a common temporal mode $\lambda(t)$. The associated dynamics is straightforward: From initially well-defined boundary conditions (the source terms $h_i(x)$), the genes increase and decrease almost proportionally and simultaneously ($\lambda(t)$) (Fig. 3D). Small rotational parts of the flow can be further accounted for by a

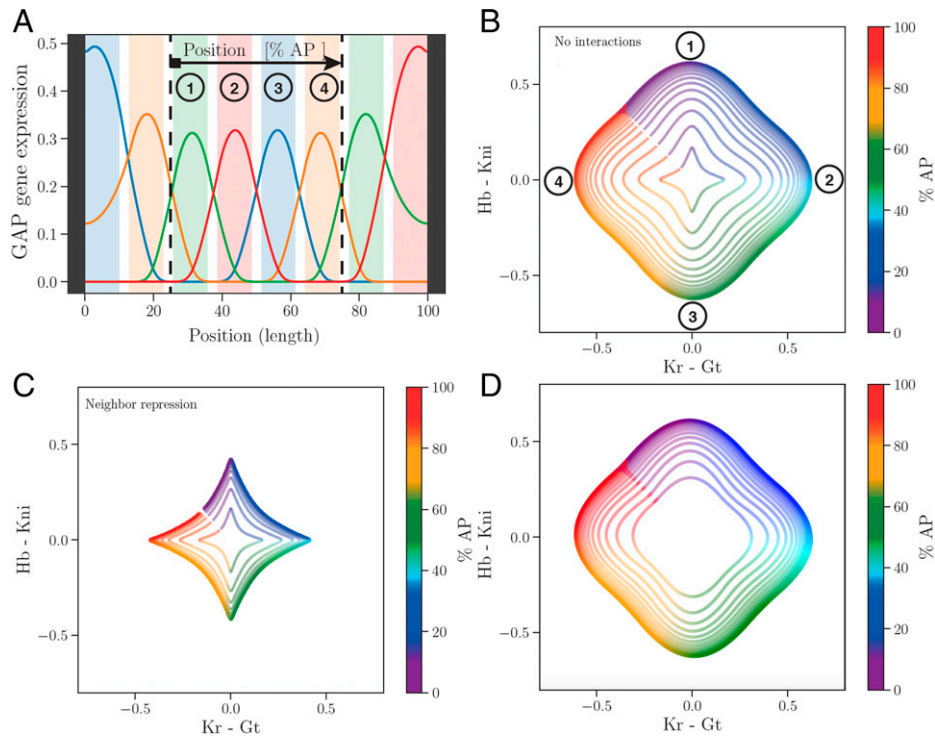


Fig. 3. Geometry of minimal models in latent space. We consider a system of PDEs with four components with localized source regions, diffusion, and degradation (SI Appendix, Eq. S7). Interactions between components are nonlinear, given by a linear interaction matrix multiplied with a Hill function to ensure positivity of the solution. More details are given in SI Appendix, section 6. (A) Stationary profile (numerical) without interactions. Source regions are shown as background color, black boxes denote the boundary of the simulation box, dashed lines denote the region shown in B–D, and numbers point out maxima. (B) Dynamic solution without interactions. Different lines show the solution at different times; smaller amplitudes in latent space correspond to earlier times. Numbers point out concentration maxima (compare to A). (C) Dynamic solution with repression of adjacent genes and next-nearest neighbors genes. (D) Dynamic solution of an “adiabatic” model, Eq. 5, showing radial expansion with almost no change of shape of the position manifold.

single local interaction (corresponding to gt and hb in the data; SI Appendix, section 6) to result in a position manifold shape in latent space that closely resembles our data (SI Appendix, Fig. S18).

Single-Mode Model with Almost No Interactions Quantitatively Recapitulates Gap Gene Dynamics. Our variable-separating toy model (Eq. 5) predicts a dominant first mode (radial flow), and all gap genes share the same temporal modes. To directly test these predictions on our data, we perform a singular value decomposition (SVD) for individual gap genes G (after removal of the shift of Fig. 2E),

$$G(x, t) = \sum_i \lambda_i^G f_i^G(t) g_i^G(x), \quad [6]$$

with the spatial ($g_i^G(x)$) and temporal ($f_i^G(t)$) parts of the dominant mode (λ_1^G) depicted in Figs. 4A and B, respectively. As predicted by the toy model, the strong first mode captures, indeed, $\lambda_1^2/(\lambda_1^2 + \lambda_2^2) \geq 97\%$ of the dynamics (SI Appendix, section 8). In addition, we see almost identical shapes of the first temporal modes for individual gap genes (Fig. 4B), again confirming the toy model prediction. From the shape of these temporal modes, we infer a two-phase process, with an initial global constant activation followed by a shutting down phase around $t \approx 25$ min (SI Appendix, section 8).

Fig. 4C and E shows the concentration fields reconstructed using only the dominant mode for each gap gene for time points 30 and 45 min, respectively (see Movie S2 for the entire dynamics). Most dynamic features of gap domains are well captured by this model, for example, position and slope of Kr and kni domains, and boundaries of gt . This is especially surprising since it is generally assumed that these features necessitate explicit mutual repression among the gap genes (17), which are not included here.

This model provides an unequivocal biological interpretation: $g_1^G(x)$ is the first spatial mode encoding the positional information provided by local factors (e.g., maternal gradients), and $f_1^G(t)$ captures a global temporal regulation. Small discrepancies of this single-mode model can be accounted for by inclusion of weak, second-order effects. For instance, addition of two linear parameters allows recovery of the weak rotational flow in latent space (corresponding to repression of gt by hb , and hb autoactivation [Fig. 4D and F, Movie S3, and SI Appendix, section 8]). Inclusion of the late second SVD mode allows recovery of the late n.c. 14 structure in the dynamics for anterior gt , hb , and Kr (SI Appendix, section 8).

Maternal Mutants Display Modularity in Latent Space. The most plausible explanation for the appearance of stationary spatial modes ($g_1^G(x)$) for each gap gene is that they are defined early during development by maternal morphogens. This interpretation provides a clear prediction: In loss-of-function maternal mutants, spatial modes should be lost where associated maternal genes are not expressed, and, in latent space, the corresponding flow should be missing. If the temporal modes are also conserved across gap genes in these maternal mutants (e.g., due to some global regulation), the flow of the remaining spatial modes in latent space should be unchanged from the flow observed in WT. Hence, gaps in embryonic development (due to the absence either of gap genes (14) or of maternal genes) should have corresponding gaps for the flow in latent space.

Gap gene expression data for maternal mutants (32) qualitatively confirm these predictions (Fig. 5). Loss of functions for maternal mutants for bcd , nos , and tor display gaps in the latent space region where the missing maternal genes are normally expressed. The flow outside those regions is qualitatively preserved

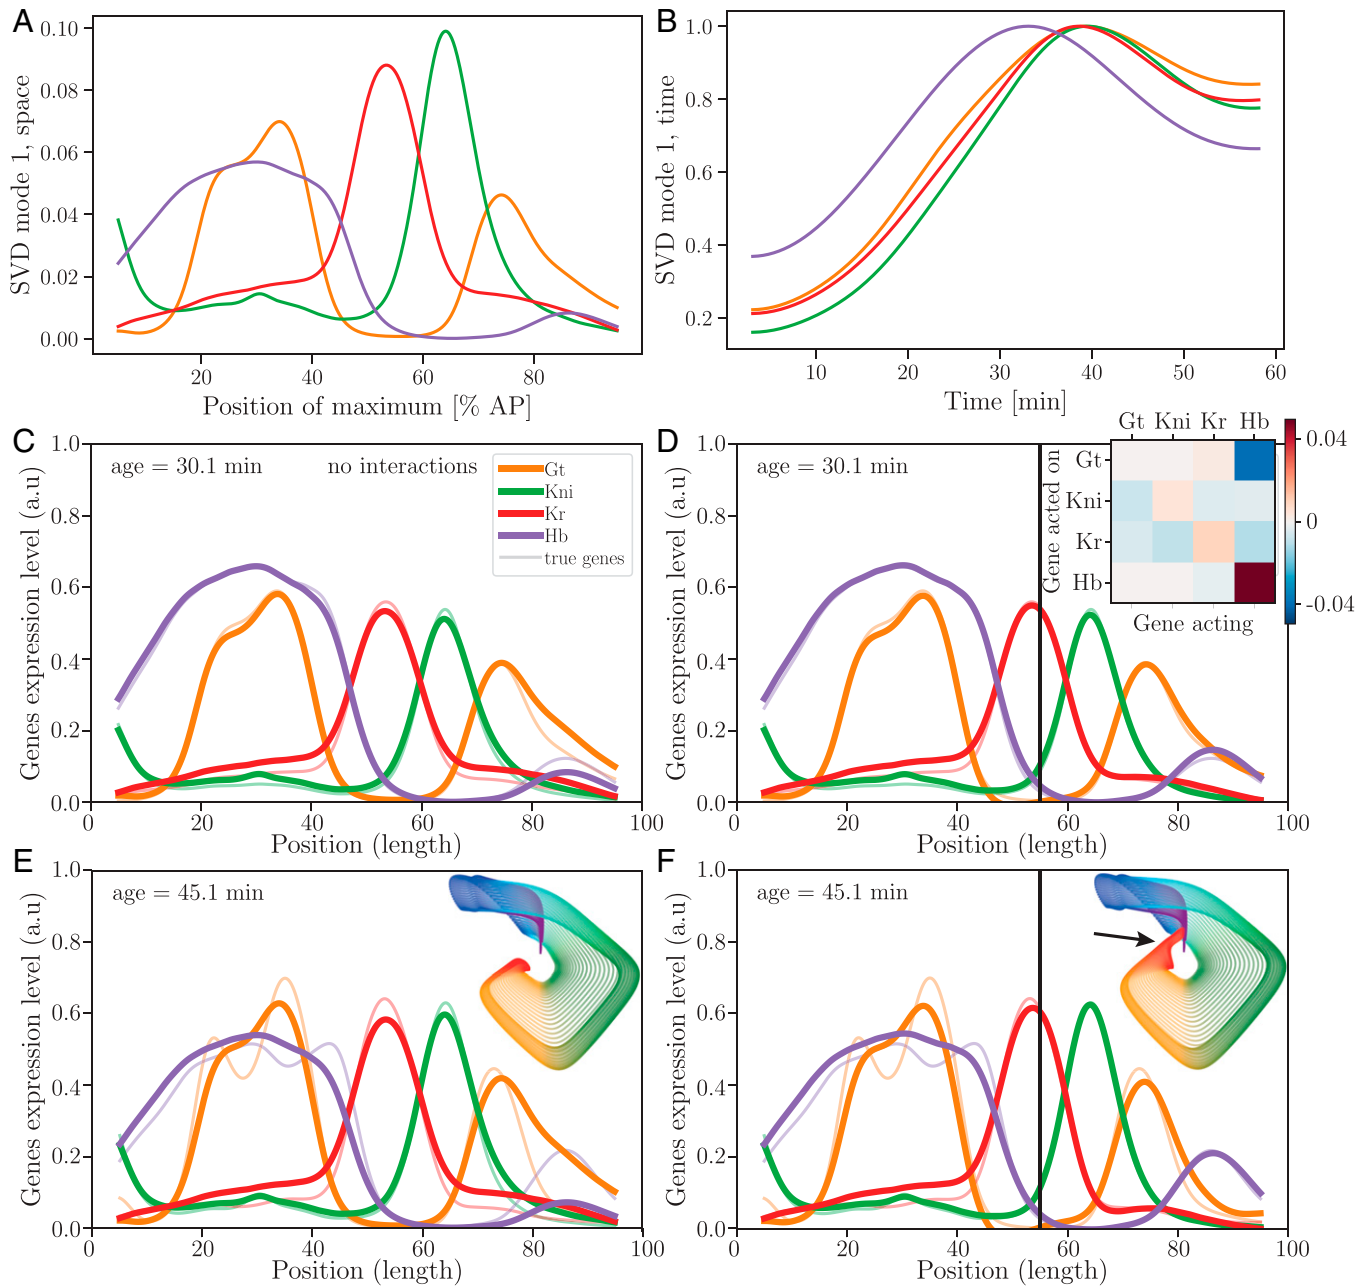


Fig. 4. Reconstruction of gap gene dynamics using first mode of SVD and posterior interactions. (A) SVD mode 1, spatial part, of the different genes (lines), colors are similar to Fig. 1. (B) SVD mode 1, temporal part (lines). Since the absolute magnitude of the modes are arbitrary, all curves are rescaled to have the same maximum value to best compare their shape. (C) Reconstruction of concentration profiles using the first mode SVD (shown for $t = 30$ min) without any interaction. (D) Reconstruction using the first mode SVD and a fit of interactions in the posterior (for positions $>55\%$ AP). The interactions are shown in *Inset* (see mathematical details in *SI Appendix, section 8*). (E and F) Same as C and D, but for $t = 45$ min. *Insets* show the dynamics in latent space.

from the WT flow (Fig. 5A). The *bcd* mutants (Fig. 5B) are restricted to the bottom left part of the latent space, *nos* mutants (Fig. 5C) are restricted to the upper part of the latent space, and *tor* mutants (Fig. 5D) truncate the position manifold at landmark B (Fig. 2A).

Double mutants (Fig. 5E–G) essentially confirm these observations, with further truncation of the flow. More quantitatively, mutant spatial modes are either truncated or stretched versions of the corresponding WT ones (Fig. 5H). For instance, the spatial modes of *nos* mutants are virtually identical to WT in the anterior (up to a multiplicative constant accounting for relative peak magnitudes), but are essentially flat in the posterior. Other mutants show similar behavior (*SI Appendix, section 8*). Temporal modes for maternal mutants are identical to WT when

bcd is present, and are delayed otherwise (*SI Appendix, Fig. S21*), possibly indicative of differential temporal regulations by maternal genes (43).

We also explore the flow in latent space for these mutants for potential gap gene interactions that might be otherwise hidden in the WT data. Again, rotational flow in latent space (similar to what we see in quadrant 2 for WT data) would be a clear signature. For the single-copy *bcd* mutant, we see weak clockwise rotational flow in quadrant 1, possibly suggestive of repression of *hb* by *Kr* (*SI Appendix, Figs. S4 and S29*). For *nos* mutants, we see weak counterclockwise rotation in quadrant 1 (*SI Appendix, Figs. S4 and S29*), suggestive of repression of *Kr* by *hb*. In both cases, however, the effects are rather mild, again pointing to a minor role for repressive interactions in this system.

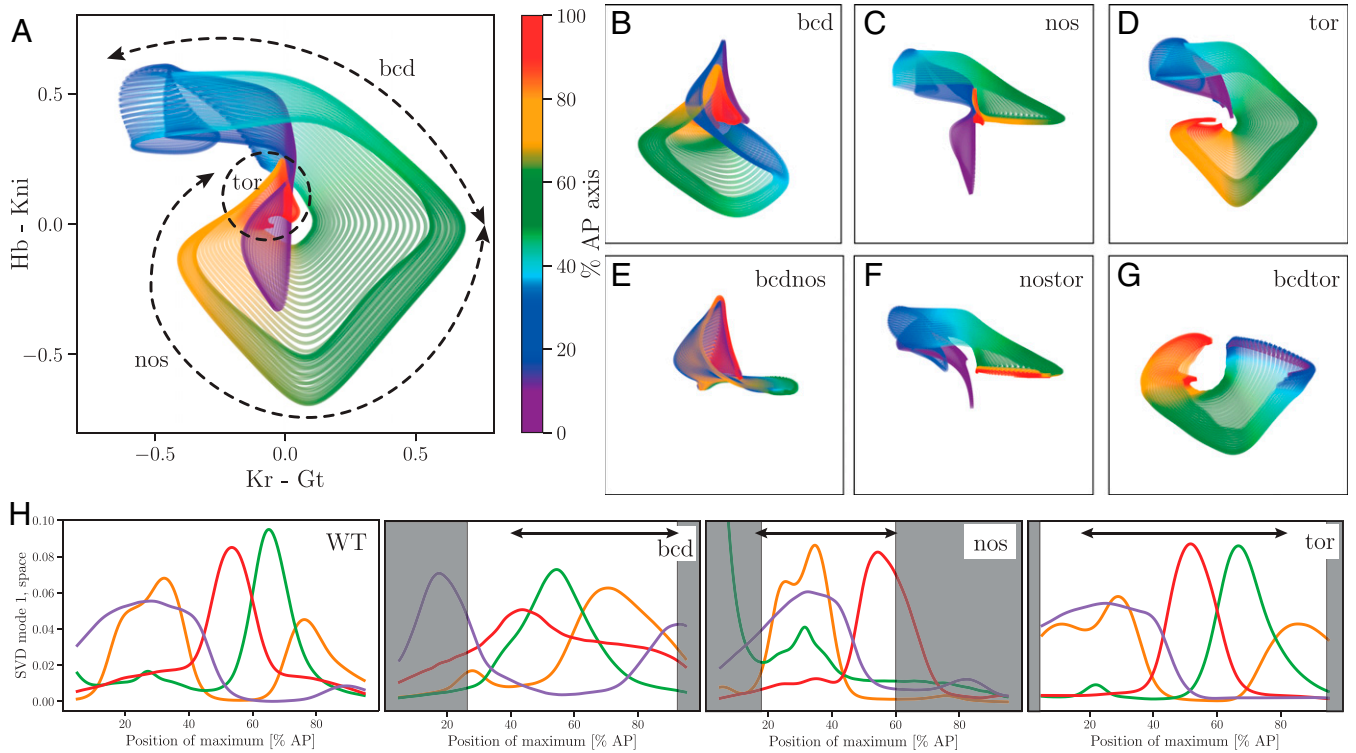


Fig. 5. Comparison of WT and maternal mutants in latent space. (A) WT with sketch of latent space angle influenced by the maternal genes *bicoid* (*bcd*), *nanos* (*nos*), and *torso* (*tor*). (B–G) Maternal mutants in latent space with deletions of one (B–D) and two (E–G) maternal genes. The names correspond to the deleted genes and miss the regions influenced by the respective genes (compare with A). (H) First spatial SVD modes of gap genes for WT and single gene mutants. Single mutants have regions comparable to the WT (white background); the comparable positions of the WT are shown as arrows.

Positional Information in Latent Space. If maternal genes define independent spatial and temporal modes, most positional information should be present early on and might be visible in latent space. Following the same procedure as Petkova et al. (32), we thus quantify the positional information carried by the position manifold. We build a 2D Bayesian map $P(x|H_1, H_2)$ from individual 40- to 44-min-old embryos to infer positions from the two latent space variables H_1 and H_2 as defined by Eqs. 1 and 2, leading to the associated position manifold in Fig. 6A (see *SI Appendix, section 2* for details). The associated decoding map (Fig. 6B) appears almost identical to the map based on the four gap genes directly (32) (Fig. 6B, *Inset*). We can reconstruct position with a precision of around 1% (*SI Appendix, section 2*), similar to the original decoding map (32), confirming that positional information is well encoded in latent space.

Similar to the original analysis (32), our position manifold-based decoding produces a region with reduced precision in the 20 to 40% egg length region, where a symmetry in the probability structure appears. This region corresponds to a singularity in the corner of the position manifold, and, physically, to the anterior region of the embryo where a small *kni* peak is missed by the autoencoder.

The big advantage of the latent space representation is that it provides a geometrical interpretation for positional information that maps gap gene concentrations to position. When we plot the most probable position given by the 2D Bayesian map (Fig. 6C), the color map defined by position x largely recapitulates the polar angle φ in latent space (at least in the 35 to 75% egg length region of the embryo, which almost entirely covers the pair-rule gene *Eve* expression domain [Fig. 6D]). From the definition of H_1 and H_2 (Eqs. 1–4), this polar angle φ corresponds to the concentration ratio of adjacent gap genes: for the first quadrant ($\varphi \in [0, \pi/2]$), $\tan \varphi = hb/Kr$; for the second

quadrant ($\varphi \in [\pi/2, \pi]$), $\tan(\pi - \varphi) = hb/gt$; for the third quadrant ($\varphi \in [\pi, 3/2\pi]$), $\tan \varphi = kni/gt$; and, for the fourth quadrant ($\varphi \in [3/2\pi, 2\pi]$), $\tan(-\varphi) = kni/Kr$. More explicitly, at each transition between two quadrants, the concentration of one of the gap genes approaches zero; for instance, the value $\varphi = 3\pi/2$ corresponds to where $H_1 \simeq gt \simeq Kr \simeq 0$, so that both ratios *kni/gt* and *kni/Kr* are singular, and $|\tan \varphi| = \infty$.

Repeating this analysis for the various mutant datasets from above (Fig. 6A), we reconstruct and recover the mutant decoding maps from Petkova et al. (32). We see strong agreement between positional information encoded either by H_1 and H_2 or by the four gap genes directly. For instance, Fig. 6E and F shows the position manifold and inferred position, respectively, for a *nos* mutant using the WT map, (map from ref. 32 reproduced in Fig. 6F, *Inset*). More examples are given in *SI Appendix, section 2*, again with excellent agreement between the two approaches. Thus the reconstructed 2D latent space from WT data transparently contains most of positional information encoded by gap genes, confirming its relevance as an analytical tool.

Discussion

We have applied a dimensionality reducing autoencoder to the *Drosophila* gap gene network and show that many features of this complex spatiotemporal system can be inferred from an unbiased 2D latent space projection. Gap gene expression dynamics is concisely described in this latent space by a simple two-mode model, one for space and one for time, and the latter is common to all gap genes, suggesting common global transcriptional or translational regulation. The entire WT dynamics can thus be explained, with few added interactions, with the spatial modes being defined during early stages by the upstream maternal genes.

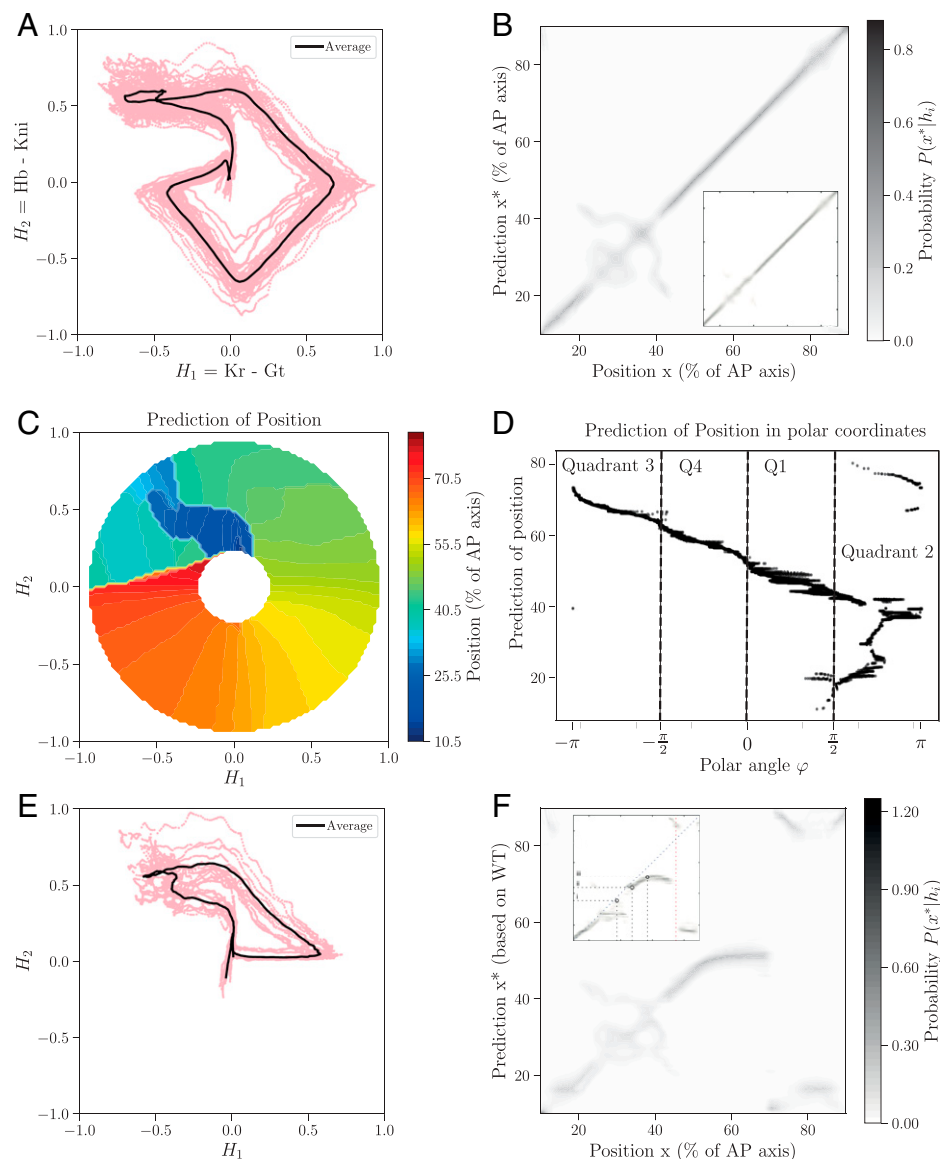


Fig. 6. Encoding positional information in latent space. (A) Overlay of position manifolds for 23 individual WT embryos at 40 min to 44 min after mitosis 13 in latent space. The average across these embryos is shown in black. (B) Bayesian probability of being at position x^* as a function of actual position x , evaluated using the latent space variables H_1 and H_2 . (Inset) Similar map evaluated with four gap genes reproduced from ref. 32. (C) Most probable position along the AP axis in latent space (shown on a circular crown to highlight correlation to polar angle). (D) Positions as a function of polar angles for all points of the circular crown from C. (E and F) Data in latent space and prediction of position for oskar mutant. (Inset) Similar map evaluated with four gap genes reproduced from ref. 32, with permission from Elsevier.

Consistent with this model, maternal gene mutants can be projected in the same latent space, validating the model's modular structure where different parts of the latent space are indeed controlled by different maternal morphogens (Fig. 5A). These results are consistent with earlier observations that maternal mutants cannot produce novel cellular fates (44), and the latent space representation demonstrates that the overall expression dynamics in these maternal mutants is largely preserved.

It has been shown previously that all four gap genes are needed to decode cellular identities, and thus position along the AP axis, with 1% accuracy (32). We have shown here that positional information can be reconstructed from the 2D latent space map, also with 1% accuracy. It might thus seem counterintuitive that two variables can achieve the same accuracy as all four gap genes together in the full system. However, by design of the autoencoder, one can reconstruct the four gap gene concentrations from the two variables H_1 and H_2 , and thus the overall information content

is conserved in the 2D latent space. A big advantage of this 2D representation is that it allows for an intuitive geometric representation of the positional information as a polar angle φ . The simplicity of the autoencoder further allows connecting this angle back to the ratio between spatially adjacent gap genes, revealing the direct connection between gap gene concentrations and positional information.

This encoding of position by the ratio of adjacent gap gene concentrations is thus directly biologically interpretable. In particular, the positions of the downstream *Eve* stripes are regulated by the two axes learned by the autoencoder. Stripes 2 and 5 are both regulated by *Kr* and *gt* (45). They are expressed where the latent variable $H_2 = Kr - Gt \approx 0$ and where *Hb* and *Kni* peak (32), corresponding to $\varphi \approx \pm\pi/2$. The stripe pairs 3/7 and 4/6 are regulated by *hb* and *kni* (26, 46); in particular, stripe 4 is expressed in a region where *kni* and *Kr* have comparable concentrations. We would then expect the enhancer for *eve*-stripe-4

to effectively compute the ratio Kr/kni , a prediction that can be checked experimentally by looking at homozygous mutants for both genes. Connecting position this way, directly to gap gene concentration ratios, provides an intuitive rationale for the modular structure of the gap gene system where maternal genes regulate those ratios by coordinating the spatial gap genes modes.

Dimensionality reduction leads to a more global view of the underlying dynamics of the data, revealing common features invisible to traditional network-based mechanistic descriptions. The latent space projects time and position onto polar coordinates radius and angle, respectively, corresponding to different SVD modes, and many dynamical features of the system become clearly visible in latent space. Position is encoded by a single spatial mode for each individual gap gene, defining distinct quadrants in latent space. A common temporal mode of all gap genes suggests a global coordination of transcriptional inputs, giving rise to uniform expansion of the position manifold in latent space. Declarations of these modes likely happen early, most probably by maternal inputs, with few modulations by downstream interactions. Latent space further reveals the underlying connection between gap gene dynamics and positional information: The dynamics mostly occurs along lines of constant polar angles φ , associated with gap gene ratios, which further defines position.

The underlying simplicity of the gap gene system uncovered by our approach calls into question the typically assumed necessity for multiple interactions that have been described in the literature (17). As mentioned above, it is generally assumed that gap genes position themselves via mutual repression, thereby generating sets of spatially adjacent gap gene expression domains. Capturing such nonlinearities was a primary motivation for the use of an autoencoder (instead of linear techniques) to understand data structure. While our autoencoder learns in an unsupervised way that the (Kr,gt) and (kni,hb) pairs are not expressed simultaneously [i.e., the “alternating cushion” model (17, 21, 27)], we also demonstrate that a single-mode (i.e., linear) model reproduces the data without explicit mutual repression in our formulation. Noteworthy is that both the location and sharpness of gap gene pattern boundaries are well explained by this model. Of course, one cannot expect a model derived from WT data to infer all existing interactions between gap genes. In particular, some existing interactions might have secondary roles, for example, for developmental robustness in response to various perturbations. But, assuming a minimal model designed to recover the spatiotemporal components of expression profiles, a strong prediction is that, in gap gene null mutants, one should not expect to see significant changes for the expression profiles of the remaining gap genes.

Existing data on gap gene expression in mutant backgrounds lack the quantitative and dynamic precision to be examined by our approach. However, in published gap gene mutant backgrounds, the changes in the expression profiles of the unperturbed gap genes appear surprisingly mild (17). In particular, in gt null mutants, the central Kr domain barely changes (47, 48), which is completely unexpected if domains are carved out by mutual repression. Similarly, in Kr null mutants, both gt domains expand toward the center but do not merge (25), meaning that other regulation is at play to carve out the gt domains (e.g., activation modulation as in our model, or some redundancies in the system). For the (kni,hb) pair, as explicitly pointed out in ref. 21, expression of hb “appears unchanged in kni null mutant embryos,” again inconsistent with an alternating cushion model with repression of hb by kni . However, in hb null mutants, the central kni domain expands strongly anteriorly (49), consistent with the alternate cushion idea. Noteworthy, however, is that hb is the only gap gene that also has a maternal component, which could play an earlier role.

More direct evidence for repression between nonoverlapping gap genes comes from gain-of-function experiments (17). For instance, overexpression of gt shifts Kr (25), but this should be contrasted with the fact that gt null mutants present no change of the central Kr domain (47, 48). As stated above, it could be that some of the interactions (here, repression of Kr by gt) mostly provide developmental robustness in a specific context (e.g., in response to overexpression of other genes), and therefore could not be inferred from WT data or even loss-of-function mutants. As illustrated by our toy model (*SI Appendix, section 6*), in the presence of relatively strong self-activation combined with local repression, mutual repressions can help stabilize gap gene kinetics and prevent further amplification (*SI Appendix, Figs. S18–S21*).

For adjacent gap gene pairs ((Kr,hb) , (Kr,kni) , (gt,kni) , and (gt,hb)), we expect any repressive interactions to be weak, since there is no mutual exclusion between those gene pairs, and, indeed, genetic evidence is sparse at best (17). In our 2D description, such repression would, in principle, be easily detectable as nonradial components of the flow. Indeed, our analysis of WT data provides a clear geometric signature for the (gt,hb) repression system in the posterior, which is consistent with experimental evidence from mutants (17, 47). However, for no other pair do we get a clear signature for repression. For example, while, for the (hb,Kr) pair, mutual repression has been suggested, possibly giving rise to multistability and canalization (42, 50), the flow we see in latent space in the corresponding region (quadrant 1 [*SI Appendix, Fig. S3* and Fig. 2E] after the very weak speed removal) is radial, and no nonlinear effects characteristic of fixed point attractors are seen (compare, e.g., with figure 7 in ref. 42).

We observed a linear speed relation for the gap gene peaks, implying a very slow anterior to posterior motion in the anterior part of the embryo and the well-described posterior to anterior flow for posterior gap genes (33). In the spirit of minimal modeling, it is more parsimonious to assume that a single, global, underlying mechanism generates the slow drift of the entire system (e.g., a slow modulation of the dynamics of the upstream maternal gradients), rather than fine-tuning of multiple weak local mutual repressive interactions. For instance, cross-interaction between posterior and anterior maternal gradients might explain such slow motion; see, for example, a model of motion for kink/antikink pairs (51). In addition, gap gene motion is amplified (damped) in bcd (nos) mutants (*SI Appendix, section 3*), which is consistent with the remarkable similarity and nonmonotonicity of the temporal modes for individual gap genes, directly suggesting a collective temporal coordination of gap genes by upstream signals.

Definite answers on the role and magnitude of interactions between gap genes might only come from precise quantification of gap gene mutant dynamics. In particular, if temporal modes and angular speed in latent space are unchanged compared to WT, this would rather suggest an important role for upstream controls, while major changes of temporal dynamics (e.g., anterior to posterior speed of gap genes) would rather point toward a role of gap gene cross-repression. It should also be pointed out that model inference from data is, of course, contingent on the assumed underlying mathematical model. Additional hypotheses, for example, inclusion of different enhancers activated at different times (52), or more complicated computations at the transcriptional level [memory, delay, irreversibility (53)], might allow for nontrivial dynamics explaining gap gene shifts. This also justifies starting with a parsimonious, data-driven approach to first decouple the description of the dynamics before assuming any underlying model. Here, we find it informative that a very simple dynamics emerges and can thus be explained without the need of multiple complex interactions.

Data Availability. Previously published data were used for this work (see ref. 32). Data are available at the following public link: [https://www.cell.com/cell/fulltext/S0092-8674\(19\)30040-6?returnURL=https%3A%2F%2Flinkinghub.elsevier.com%2Fretrieve%2Fpii%2FS0092867419300406%3Fshowall%3Dtrue#supplementaryMaterial](https://www.cell.com/cell/fulltext/S0092-8674(19)30040-6?returnURL=https%3A%2F%2Flinkinghub.elsevier.com%2Fretrieve%2Fpii%2FS0092867419300406%3Fshowall%3Dtrue#supplementaryMaterial).

ACKNOWLEDGMENTS. We thank Marianne Bauer, William Bialek, and Eric Wieschaus for discussion and comments. This work was supported, in part, by the US NSF, through the Center for the Physics of Biological Function (Grant PHY-1734030); by NIH Grants R01GM097275, U01DA047730, and U01DK127429;

1. A. D. Lander, The edges of understanding. *BMC Biol.* **8**, 40 (2010).
2. J. Gunawardena, Models in biology: 'Accurate descriptions of our pathetic thinking.' *BMC Biol.* **12**, 29 (2014).
3. W. Bialek, *Biophysics: Searching for Principles* (Princeton University Press, 2012).
4. I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning* (MIT Press, 2016).
5. S. M. Udrescu, M. Tegmark, Al Feynman: A physics-inspired method for symbolic regression. *Sci. Adv.* **6**, eaay2631 (2020).
6. F. Corson, E. D. Siggia, Geometry, epistasis, and developmental patterning. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 5568–5575 (2012).
7. J. A. Fritz *et al.*, Shared developmental programme strongly constrains beak shape diversity in songbirds. *Nat. Commun.* **5**, 3700 (2014).
8. F. Corson, L. Couturier, H. Rouault, K. Mazouni, F. Schweisguth, Self-organized Notch dynamics generate stereotyped sensory organ patterns in *Drosophila*. *Science* **356**, eaai7407 (2017).
9. F. Corson, E. D. Siggia, Gene-free methodology for cell fate dynamics during development. *eLife* **6**, e30743 (2017).
10. F. Proulx-Girardeau, T. J. Rademaker, P. François, Untangling the hairball: Fitness-based asymptotic reduction of biological networks. *Biophys. J.* **113**, 1893–1906 (2017).
11. L. Jutras-Dubé, E. El-Sherif, P. François, Geometric models for robust encoding of dynamical information into embryonic patterns. *eLife* **9**, e55778 (2020).
12. M. Saez *et al.*, Statistically derived geometrical landscapes capture principles of decision-making dynamics during cell fate transitions. *Cell Syst.* **13**, 1 (2022).
13. B. C. Daniels, I. Nemenman, Automated adaptive inference of phenomenological dynamical models. *Nat. Commun.* **6**, 8133 (2015).
14. C. Nüsslein-Volhard, E. Wieschaus, Mutations affecting segment number and polarity in *Drosophila*. *Nature* **287**, 795–801 (1980).
15. J. P. Gergen, D. Coulter, E. F. Wieschaus, Segmental pattern and blastoderm cell identities. *Symp. Soc. Dev. Biol.* **43**, 195–220 (1986).
16. P. Lawrence, *The Making of a Fly: The Genetics of Animal Design* (Blackwell Scientific, Oxford, 1992).
17. J. Jaeger, The gap gene network. *Cell. Mol. Life Sci.* **68**, 243–274 (2011).
18. J. Briscoe, S. Small, Morphogen rules: Design principles of gradient-mediated embryo patterning. *Development* **142**, 3996–4009 (2015).
19. H. Meinhardt, Hierarchical inductions of cell states: A model for segmentation in *Drosophila*. *J. Cell Sci. Suppl.* **4**, 357–381 (1986).
20. S. B. Carroll, S. H. Vavra, The zygotic control of *Drosophila* pair-rule gene expression. II. Spatial repression by gap and pair-rule gene products. *Development* **107**, 673–683 (1989).
21. H. Jäckle, D. Tautz, R. Schuh, E. Seifert, R. Lehmann, Cross-regulatory interactions among the gap genes of *Drosophila*. *Nature* **324**, 668–670 (1986).
22. J. Treisman, C. Desplan, The products of the *Drosophila* gap genes *hunchback* and *Krüppel* bind to the *hunchback* promoters. *Nature* **341**, 335–337 (1989).
23. M. Hoch, E. Seifert, H. Jäckle, Gene expression mediated by *cis*-acting sequences of the *Krüppel* gene in response to the *Drosophila* morphogens *bicoid* and *hunchback*. *EMBO J.* **10**, 2267–2278 (1991).
24. J. O. Dubuis, G. Tkacik, E. F. Wieschaus, T. Gregor, W. Bialek, Positional information, in bits. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 16301–16308 (2013).
25. R. Kraut, M. Levine, Mutually repressive interactions between the gap genes giant and *Krüppel* define middle body regions of the *Drosophila* embryo. *Development* **111**, 611–621 (1991).
26. D. E. Clyde *et al.*, A self-organizing system of repressor gradients establishes segmental complexity in *Drosophila*. *Nature* **426**, 849–853 (2003).
27. D. Papatsenko, M. Levine, The *Drosophila* gap gene network is composed of two parallel toggle switches. *PLoS One* **6**, e21145 (2011).
28. R. Rivera-Pomar, H. Jäckle, From gradients to stripes in *Drosophila* embryogenesis: Filling in the gaps. *Trends Genet.* **12**, 478–483 (1996).
29. P. A. Lawrence, Background to bicoid. *Cell* **54**, 1–2 (1988).
30. L. Wolpert, *Principles of Development* (Oxford University Press, 2006).
31. G. Tkačik, T. Gregor, The many bits of positional information. *Development* **148**, dev176065 (2021).
32. M. D. Petkova, G. Tkačik, W. Bialek, E. F. Wieschaus, T. Gregor, Optimal decoding of cellular identities in a genetic network. *Cell* **176**, 844–855.e15 (2019).
33. J. Jaeger *et al.*, Dynamical analysis of regulatory interactions in the gap gene system of *Drosophila melanogaster*. *Genetics* **167**, 1721–1737 (2004).
34. A. Crombach, K. R. Wotton, E. Jiménez-Guri, J. Jaeger, Gap gene regulatory dynamics evolve along a genotype network. *Mol. Biol. Evol.* **33**, 1293–1307 (2016).
35. B. Verd, N. A. Monk, J. Jaeger, Modularity, criticality, and evolvability of a developmental gene regulatory network. *eLife* **8**, e42832 (2019).
36. G. E. Hinton, R. R. Salakhutdinov, Reducing the dimensionality of data with neural networks. *Science* **313**, 504–507 (2006).
37. M. Bauer, M. D. Petkova, T. Gregor, E. F. Wieschaus, W. Bialek, Trading bits in the readout from a genetic network. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2109011118 (2021).
38. L. v. d. Maaten, G. Hinton, Visualizing data using T-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
39. L. McInnes, J. Healy, J. Melville, UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv [Preprint]* (2018). <https://doi.org/10.48550/arXiv.1802.03426>. Accessed 17 August 2020.
40. J. Jaeger *et al.*, Dynamic control of positional information in the early *Drosophila* embryo. *Nature* **430**, 368–371 (2004).
41. D. Krotov, J. O. Dubuis, T. Gregor, W. Bialek, Morphogenesis at criticality. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 3683–3688 (2014).
42. Manu *et al.*, Canalization of gene expression in the *Drosophila* blastoderm by gap gene cross regulation. *PLoS Biol.* **7**, e1000049 (2009).
43. F. Liu, A. H. Morrison, T. Gregor, Dynamic interpretation of maternal inputs by the *Drosophila* segmentation gene network. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 6724–6729 (2013).
44. M. V. Staller *et al.*, A gene expression atlas of a *bicoid*-depleted *Drosophila* embryo reveals early canalization of cell fate. *Development* **142**, 587–596 (2015).
45. D. Stanojevic, S. Small, M. Levine, Regulation of a segmentation stripe by overlapping activators and repressors in the *Drosophila* embryo. *Science* **254**, 1385–1387 (1991).
46. M. Fujioka, Y. Emi-Sarker, G. L. Yusibova, T. Goto, J. B. Jaynes, Analysis of an even-skipped rescue transgene reveals both composite and discrete neuronal and early blastoderm enhancers, and multi-stripe positioning by gap gene repressor gradients. *Development* **126**, 2527–2538 (1999).
47. E. D. Eldon, V. Pirrotta, Interactions of the *Drosophila* gap gene giant with maternal and zygotic pattern-forming genes. *Development* **111**, 367–378 (1991).
48. U. Gaul, H. Jäckle, Pole region-dependent repression of the *Drosophila* gap gene *Krüppel* by maternal gene products. *Cell* **51**, 549–555 (1987).
49. M. Hülkamp, C. Pfeifle, D. Tautz, A morphogenetic gradient of *hunchback* protein organizes the expression of the gap genes *Krüppel* and *knirps* in the early *Drosophila* embryo. *Nature* **346**, 577–580 (1990).
50. Manu *et al.*, Canalization of gene expression and domain shifts in the *Drosophila* blastoderm by dynamical attractors. *PLoS Comput. Biol.* **5**, e1000303 (2009).
51. S. Vakulenko, Manu, J. Reinitz, O. Radulescu, Size regulation in the segmentation of *Drosophila*: Interacting interfaces between localized domains of gene expression ensure robust spatial patterning. *Phys. Rev. Lett.* **103**, 168102 (2009).
52. E. El-Sherif, M. Levine, Shadow enhancers mediate dynamic shifts of gap gene expression in the *Drosophila* embryo. *Curr. Biol.* **26**, 1164–1169 (2016).
53. J. Desponds, M. Vergassola, A. M. Walczak, A mechanism for *hunchback* promoters to readout morphogenetic positional information in less than a minute. *eLife* **9**, e49758 (2020).

by the Natural Sciences and Engineering Research Council of Canada, Discovery Grant program; and by the Simons Foundation, Mathematical Modeling of Living System program.

Author affiliations: ¹Physics Department, McGill University, Montreal, QC H3A2T8, Canada; ²Université Lyon, Ecole Normale Supérieure de Lyon, Laboratoire de Physique, UMR 5672, Université Claude Bernard Lyon 1, CNRS, 69342, Lyon, France; ³Program in Biophysics, Harvard University, Cambridge, MA 02138; ⁴Joseph Henry Laboratory of Physics, Princeton University, Princeton, NJ 08544; ⁵Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton, NJ 08544; and ⁶Department of Stem Cell and Developmental Biology, UMR3738, Institut Pasteur, 75015 Paris, France