



HAL
open science

Trading bits in the readout from a genetic network

Marianne Bauer, Mariela D Petkova, Thomas Gregor, Eric F Wieschaus,
William Bialek

► **To cite this version:**

Marianne Bauer, Mariela D Petkova, Thomas Gregor, Eric F Wieschaus, William Bialek. Trading bits in the readout from a genetic network. 2020. pasteur-03216308v1

HAL Id: pasteur-03216308

<https://pasteur.hal.science/pasteur-03216308v1>

Preprint submitted on 3 May 2021 (v1), last revised 13 Dec 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Trading bits in the readout from a genetic network

Marianne Bauer,^{1–4} Mariela D. Petkova,⁵ Thomas Gregor,^{1,2,6} Eric F. Wieschaus,^{2–4} and William Bialek^{1,2,7}

¹Joseph Henry Laboratories of Physics, ²Lewis-Sigler Institute for Integrative Genomics,

³Department of Molecular Biology, and ⁴Howard Hughes Medical Institute,
Princeton University, Princeton, NJ 08544 USA

⁵Program in Biophysics, Harvard University, Cambridge, MA 02138, USA

⁶Department of Developmental and Stem Cell Biology,
UMR3738, Institut Pasteur, 75015 Paris, France

⁷Initiative for the Theoretical Sciences, The Graduate Center,
City University of New York, 365 Fifth Ave, New York, NY 10016

In genetic networks, information of relevance to the organism is represented by the concentrations of transcription factor molecules. In order to extract this information the cell must effectively “measure” these concentrations, but there are physical limits to the precision of these measurements. We explore this trading between bits of precision in measuring concentration and bits of relevant information that can be extracted, using the gap gene network in the early fly embryo as an example. We argue that cells in the embryo can extract all the available information about their position, but only if the concentration measurements approach the physical limits to information capacity. These limits necessitate the observed proliferation of enhancer elements with sensitivities to combinations of transcription factors, but fine tuning of the parameters of these multiple enhancers is not required.

Cells control the concentrations of proteins in part by controlling the transcription of messenger RNA. This control is effected by the binding of transcription factor (TF) proteins to specific sites along the genome. Transcription factors thus can regulate the synthesis of other TFs, forming a genetic network. The dynamical variables at the nodes of this network are the TF concentrations, which represent information of relevance to the cell. What must the cell do in order to extract and use this information?

We usually think of transcription factors as controlling the level of gene expression, but we can also view the expression level as being the cell’s measurement of the TF concentration [1, 2]. As outside observers of the cell, we can measure the concentration of transcription factors with considerable accuracy [3]. However the cell’s “measurement” of TF concentration is based on the arrival of these molecules at their binding sites, and this is a noisy process, because TF concentrations are low, in the nanoMolar range [4–8]. Physical limits to the measurement of such low concentrations were first explored in the context of bacterial chemotaxis [9], but have proven to be much more general [1, 10–12]. What will be important for our discussion is not the precise values of these limits, but rather that the limits exist and are significant on the scale of biological function.

We focus on the example of the gap genes that are crucial in the early events of embryonic development in fruit flies [13, 14]. These four genes form a network with inputs from primary maternal morphogen molecules, and outputs in the striped patterns of pair-rule gene expression. These stripes are positioned with an accuracy of $\pm 1\%$ along the long (anterior–posterior) axis of the embryo, and this is the accuracy of subsequent developmental events such as the formation of the cephalic furrow

[15, 16]. The local concentrations of the gap gene proteins provide just enough information to support this level of precision [15], and the algorithm that achieves optimal readout of this positional information predicts, quantitatively, the distortions of the pair-rule stripes in mutant flies where individual maternal inputs are deleted [17].

The gap gene network offers us the chance to ask how accurately the transcription factor concentrations need to be measured to retain crucial biological information: the information that the gap genes convey about position along the anterior–posterior axis is what allows nuclei to make distinct cell fate decisions required for development. We start with more traditional views of how information is represented in TF concentrations, and then argue for a more abstract formulation of the problem.

The classical view of the gap genes is that they are expressed in domains [14]. Implicitly this suggests that fine scale variations in the concentration of these molecules is not important; rather all that matters whether expression is on or off. The quantitative version of this idea is that subsequent events are sensitive to whether expression levels are above or below a threshold, corresponding to whether a cell is inside or outside an expression domain. We know that such simple thresholding omits a lot of the information that gap gene expression levels carry about position along the anterior–posterior axis [15].

In Figure 1A and B we use the gap gene *hunchback* (Hb) to illustrate this information loss. At each point x there is an expression level g , drawn from a probability distribution $P(g|x)$; looking at many embryos we have samples out of this distribution. If cells are sensitive only to whether expression levels are above or below a threshold θ , then the variable which matters is

$$\sigma = H(g - \theta), \quad (1)$$

where H is the Heaviside step function, $H(y > 0) = 1$

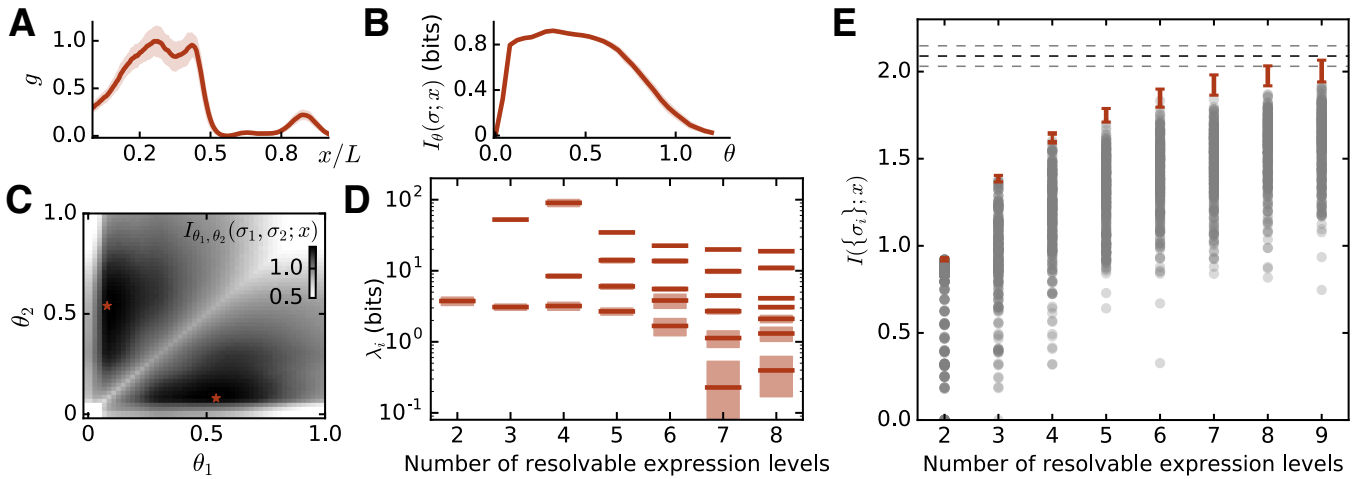


FIG. 1: (A) Hb expression level vs position along the anterior–posterior axis embryo. Mean (line) \pm one standard deviation (shading) across $N_{\text{em}} = 38$ embryos in a five minute window (40–44 min) in nuclear cycle 14 [17]. (B) Positional information vs threshold, from Eq (4). (C) Positional information with two thresholds, $I_{\theta_1, \theta_2}(\sigma_1, \sigma_2; x)$ (bits). (D) Eigenvalues $\{\lambda_i\}$ of the Hessian matrix χ , from Eq (8). The number of eigenvalues is the number of thresholds, one less than the number of resolvable expression levels. Shaded bands are \pm one standard error in our estimates. (E) Positional information captured with $i = 1, \dots, K$ thresholds, as a function of the number of resolvable levels $K + 1$. Error bars (red) are mean \pm one standard error of our estimate of the maximum. Circles (grey) are 300 values of $I(\{\sigma_i\}; x)$ at random settings of the K thresholds $\{\theta_i\}$. The black dashed line indicates the positional information $I(g; x)$, available from the expression level if measured precisely, and gray dashed lines are \pm one standard error in our estimate of this information.

and $H(y < 0) = 0$. Then we can estimate the distribution $P_\theta(\sigma|x)$, where the subscript reminds us that the result depends on our choice of threshold,

$$P_\theta(\sigma = 1|x) = \int dg H(g - \theta)P(g|x) \quad (2)$$

$$P_\theta(\sigma = 0|x) = 1 - P_\theta(\sigma = 1|x). \quad (3)$$

Finally we compute the amount of information that the discrete variable σ provides about position,

$$I_\theta(\sigma; x) = \sum_\sigma \int dx P(x)P_\theta(\sigma|x) \log_2 \left[\frac{P_\theta(\sigma|x)}{P_\theta(\sigma)} \right] \text{ bits}, \quad (4)$$

where a priori all positions along the length of the embryo are equally likely, $P(x) = 1/L$, and

$$P_\theta(\sigma) = \int dx P(x)P_\theta(\sigma|x). \quad (5)$$

It is important that in exploring the impact of thresholding we allow for the best possible choice of the threshold θ , which in this example proves to be at $\theta^* \sim 1/3$ of the maximal mean expression level.

If the expression level is represented only by the on/off or binary variable σ , then it can provide at most one bit of information (about anything), and we see that the information about position comes close to this bound, with $I_{\text{max}}(\sigma; x) = 0.92 \pm 0.01$ bits. But this is less than one half of the information that is carried by the continuous expression level, $I(g; x) = 2.09 \pm 0.06$ bits.

One path to recovering the information that was lost by thresholding is to imagine that the cell can resolve more details, perhaps distinguishing reliably among three or four different expression levels rather than just two. This is equivalent to imagining that we have multiple readout mechanisms, each of which can only distinguish on/off, but the different mechanisms can have different thresholds, in the spirit of the “French flag” model [18]. Because we can always put the thresholds in order, having K binary switches is the same as distinguishing $K + 1$ different expression levels. It can be useful to think of thresholding as being implemented at individual binding sites for the TFs, or perhaps at cooperative arrays of binding sites in enhancers, but it is important that our arguments are independent of these microscopic details.

If we have two different elements, each of which reports on whether the expression level is above or below a threshold, then the relevant variables are

$$\sigma_1 = H(g - \theta_1) \quad (6)$$

$$\sigma_2 = H(g - \theta_2). \quad (7)$$

We see in Fig 1C that there is a broad optimum in the positional information that these variables capture, $I_{\theta_1, \theta_2}(\{\sigma_1, \sigma_2\}; x)$, when the two thresholds are quite different from one another, $\theta_1^* = 0.1$ and $\theta_2^* = 0.58$; these bracket the optimal single threshold $\theta = 0.34$. The maximum information now is $I_{\text{max}}(\{\sigma_1, \sigma_2\}; x) = 1.4 \pm 0.015$ bits, noticeably more than in the case with one threshold but still far from capturing all the available information.

We can generalize this idea to multiple thresholding elements, which are described by a set of variables $\{\sigma_i\}$, with each $\sigma_i = H(g - \theta_i)$, for $i = 1, 2, \dots, K$; the relevant quantity now is $I(\{\sigma_i\}; x)$. This positional information depends on all the thresholds $\{\theta_i\}$, and we perform a multidimensional optimization to find the maximum of $I(\{\sigma_i\}; x)$. Figure 1E shows that for cells to extract all the positional information available from the Hb concentration, they must distinguish eight or nine different expression levels, in effect representing the expression level with ~ 3 bits of precision.

Distinguishing eight levels in this simple threshold picture requires the cell to set seven thresholds. It might seem as though this necessitates setting each threshold to its optimal value, a form of fine tuning. To explore this we choose thresholds at random, uniformly in the relevant interval $0 < \theta < 1$. As shown in Fig 1E, the typical random choices are far below the optimum, as expected. But Figures 1B and C show that there is broad plateau in information vs one or two thresholds, and even with eight thresholds we find that more than 1 in 1000 of our random choices come within error bars of the optimum.

Another way of looking at the issue of fine tuning is to examine the behavior of the information in the neighborhood of the optimum,

$$I(\{\theta_i\}) = I_{\max}(K) + \frac{1}{2} \sum_{i,j=1}^K (\theta_i - \theta_i^*) \chi_{ij} (\theta_j - \theta_j^*) + \dots, \quad (8)$$

estimating the Hessian matrix χ numerically from the data. The matrix χ has units of bits, as we chose the thresholds to be dimensionless. The eigenvectors of χ determine the combinations of thresholds that have independent effects on the information, and the eigenvalues $\{\lambda_i\}$ of χ (also in bits) determine the sensitivity along these independent directions. As the number of thresholds increases we find a wider spread of eigenvalues, as in a wide class of “sloppy models” [19, 20]. This means that some combinations of thresholds are two orders of magnitude more important than others.

In more detail we find that the eigenvector with the largest eigenvalue is concentrated on the lowest threshold, while higher thresholds are much more loosely constrained. Although we should again be cautious about overly detailed molecular interpretations, it is natural to think of the mapping $g \rightarrow \{\sigma_i\}$ as being implemented by binding of the transcription factor to specific sites along the genome, so that thresholds are set by the binding constants or affinities of the TF for these sites. The spectrum of the Hessian tells us that the affinity at the strongest binding site must be set carefully, but the weaker binding sites can be scattered more freely across the available dynamic range of concentrations. A near optimal array of thresholds could thus evolve by duplication of a strong binding site, followed by sequence drift to weaker binding, and then selection for the more complex and re-

producible patterns that result from capturing more positional information, as in Ref [21].

The idea of the cell taking thresholds is that it is not possible to make arbitrarily fine distinctions, so we imagine that only “on” and “off” are reliably distinguishable. But thresholding doesn’t quite capture our intuition about limited precision because the threshold itself is sharp: taking a threshold means making perfect distinctions between concentrations that are just above and just below threshold. We need to find a more precise way of implementing the intuition that cells cannot make arbitrarily precise measurements.

One path to a more realistic view of limited precision is to explore microscopic models for the interactions between TF proteins and their binding targets along the genome, and for the mechanisms by which binding events influence transcription [22–27]. But then our conclusions depend on our choice of model, and almost certainly our model will be wrong in detail. Instead we can try to abstract away from concrete models, hoping to capture the essence of limited precision in a more general way.

As an example, suppose that transcription factors influence transcription only when they are bound to their target sites. In this case the cell does not have direct access to the concentration g but only to the binding site occupancy averaged over time; there could also be multiple binding sites. Abstractly, the concentration g is being mapped into some other variable, and we can call this variable C , without worrying about details. The crucial consequence of noise in the binding events is that C can provide only a limited amount of information about the concentration g [2]. Different molecular mechanisms generate different mappings $g \rightarrow C$, but in all mechanisms the low concentrations of the relevant molecules limit the information that is transmitted, $I(C; g)$.

What is relevant for the embryo is information about the position of cells along the anterior–posterior axis. Thus it is natural to ask: given physical limits on $I(C; g)$, what is the maximum amount of positional information, $I(C; x)$, that can be captured? Mathematically we want to search over all (noisy) mappings $g \rightarrow C$, holding $I(C; g)$ fixed and maximizing $I(C; x)$:

$$\max_{P(C|g)} [I(C; x) - TI(C; g)], \quad (9)$$

where T is a Lagrange multiplier that we introduce to fix $I(C; g)$. This is the “information bottleneck” problem [29]. Buried inside the function that we are trying to maximize is the joint distribution of positions and expression levels, $P(g, x)$, which has been estimated experimentally [17], and it is this distribution that determines the form of the optimal $P(C|g)$.

We solve the optimization problem in Eq (9) numerically, considering C to be a discrete set of states and varying the number of these states $\|C\|$. At fixed $\|C\|$, decreasing T allows $I(C; g)$ to be larger, and pushes the




FIG. 2: The information bottleneck for positional information carried by Hb expression levels. We map expression into some compressed description, $g \rightarrow C$, and find the maximum $I(C;x)$ at fixed $I(C;g)$, from Eq (9), shown as the solid line with different greyshades indicating different numbers of states $\|C\|$. Solid points with error bars are the $T \rightarrow 0$ limit with fixed $\|C\|$, which reproduces the optimal discretization by multiple thresholds in Fig 1E, and light points are from an explicitly deterministic formulation of the bottleneck problem [30]. Upper panel shows snapshots probability distributions $P(C|g)$ at different information capacities along the bottleneck curve; intermediate levels of $g \in [0.05, 0.8]$ are progressively better resolved as the capacity increases.

noisy mapping $P(C|g)$ toward being deterministic. Results of the bottleneck analysis for Hb are shown in Fig 2 as trajectories in the plane $I(C;x)$ vs $I(C;g)$. The bounding curve that emerges from this analysis at large $\|C\|$ separates the plane into a physically possible region (below the curve) and an impossible region (above the curve). As $I(C;g)$ becomes large, the curve plateaus at the available positional information $I(g;x)$, and must also stay below the diagonal, so that $I(C;x) \leq I(C;g)$.

The optimal thresholding models from Fig 1 correspond to the endpoints of the bottleneck solutions with $\|C\|$ equal to the number of resolvable expression levels, and we see that these models are almost on the optimal bounding curve. This means that although the picture of multiple noiseless thresholds is physically wrong, it does capture the impact of limited information capacity, almost quantitatively. We can understand more about the structure of the optimal mappings $g \rightarrow C$ by looking at the distributions $P(C|g)$, shown in the top panels of Fig 2. At small $I(C;g)$ the whole ranges of g are mapped

uniformly into ranges of C , while at larger $I(C;g)$ we see the emergence of a reliably graded mapping, especially in the range bracketing half-maximal expression.

We see from Fig 2 that capturing all the positional information encoded by Hb requires measuring the expression level with ~ 3 bits of precision, consistent with our conclusions from the analysis of threshold models. We have done the same analysis for the other gap genes (Kr, Gt, and Kni), with the same conclusion. How does this compare with the information capacity of real genetic regulatory elements? Estimates based both on direct measurements and on more detailed models indicate that this capacity is in the range of 1–3 bits [2, 28]. This shows that it is possible for cells to extract all the available positional information using a handful of enhancers, but this requires that they operate close to the physical limits to information capacity. At the same time, it is clear from Fig 2 that there is a very big difference between a capacity of 1 bit and 3 bits.

The actual information capacity of genetic regulatory elements depends on the absolute concentrations of the relevant molecules, on the time available for reading out the information, and on other details of the different noise sources in the system [2, 28]. At one extreme, if the capacity is three bits, then a single regulatory element is sufficient to capture the full positional information. But if the capacity of a single element is only one bit, then we need multiple regulatory elements even in response to a single transcription factor.

It might seem that if each element can capture ~ 1 bit and the cell needs ~ 3 bits, then there must be three independent regulatory elements, but the problem is more




FIG. 3: Decomposing a two bit element into one bit elements. (A) Mapping from g into four levels of C , capturing $I(C;g) = 1.9$ bits. (B) Decomposing $C = (C_1, C_2)$ where each C_i has two levels and captures $I(C_i;g) \leq 1$ bit. This decomposition is exact, but the mapping $g \rightarrow C_1$ is non-monotonic. (C) Decomposing $C = (C_1, C_2, C_3)$ where each C_i has two levels and captures $I(C_i;g) \leq 1$ bit. Adding a third element allows each mapping $g \rightarrow C_i$ to be monotonic.

subtle than this. Three elements suffice if they provide independent information about the expression level g , but typically they carry overlapping or redundant information [31]. The only way to avoid this redundancy is for the mapping $g \rightarrow C$ be non-monotonic, as in Fig 3. If we try to build a molecular implementation of non-monotonicity, we will end up with the transcription factor functioning (for example) as an activator at low concentrations and a repressor at high concentrations. There is older work suggesting that gap genes can act non-monotonically on their targets [32], and recent work motivated by bacterial systems showing how more complex patterns of TF binding and DNA looping can lead to non-monotonicity [33]. There also is evidence that Hb in particular can act as an activator at some enhancer sites and a repressor at others [34].

If the cell does not have access to molecular mechanisms that allow for non-monotonicity, then there must be more of the ~ 1 bit elements. In the limit, we return to the picture of many elements, each of which switches at a different threshold, as discussed above.

Thus far our discussion has focused on the information carried by the expression level of just one gene, but we know that positional information is carried by the combined expression levels of all four gap genes [15, 17]. Rather than considering, as above, the mapping $g_{\text{Hb}} \rightarrow C$, we can consider mappings from combinations of expression levels into compressed variables. We start with Hb and Kr, so we consider $\{g_{\text{Hb}}, g_{\text{Kr}}\} \rightarrow C$. The analog of Eq (9) is the optimization problem

$$\max_{P(C|\{g_{\text{Hb}}, g_{\text{Kr}}\})} [I(C; x) - TI(C; \{g_{\text{Hb}}, g_{\text{Kr}}\})]. \quad (10)$$

Aspects of the solution are illustrated in Fig 4.

The structure of the information bottleneck curve for two genes (Fig 4B) is very similar to that for one gene. To be sure that the positional information has saturated, we need to capture ~ 6 bits of information about the combined expression levels of Hb and Kr, or ~ 3 bits per gene. Importantly, the optimal compressions cannot be decomposed into independent compressions of Hb and Kr. With ~ 6 bits in total, the mapping $\{g_{\text{Hb}}, g_{\text{Kr}}\} \rightarrow C$ divides the plane of expression levels into $\|C\| \sim 2^6 = 64$ different regions, which is complicated to visualize. For purposes of illustration we look at the optimal solution with $\|C\| = 9$, shown in Fig 4C.

The compressed state that covers the largest area in the $\{g_{\text{Hb}}, g_{\text{Kr}}\}$ plane ($C = 9$) has borders that are nearly parallel to the axes, so this is an example of a state that could be defined by thresholding the two expression levels independently. In contrast, some of the smaller areas have boundaries more nearly parallel to the diagonal, corresponding to thresholds on the ratio $g_{\text{Hb}}/g_{\text{Kr}}$. These sorts of combinatorial structures become more prominent as we capture more of information, or consider more of the gap genes simultaneously. In particular, when we in-

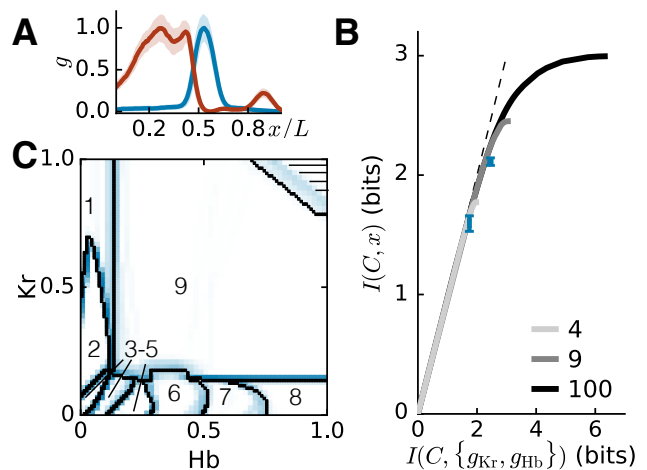


FIG. 4: The information bottleneck for positional information carried by the combination of Hb and Kr expression levels. (A) Expression vs position along the anterior-posterior axis for Hb (red) and Kr (blue). Mean (solid) and standard deviation (shading) across $N_{\text{em}} = 38$ embryos in a five minute window (40–44 min) in nuclear cycle 14 [17]. (B) Information bottleneck results, as in Fig 2, for $\|C\| \in \{4, 9, 100\}$ (shades of grey), and optimal threshold solutions with one and two thresholds for Hb and Kr each (blue); thresholded solution are worse than the optimal bottleneck solutions. (C) Structure of the optimal mapping $\{g_{\text{Hb}}, g_{\text{Kr}}\} \rightarrow C$, with $\|C\| = 9$ states in the limit $T \rightarrow 0$; solid black lines are boundaries between different states C , blue shading indicates errors in our estimates of these boundaries, and hatched region at upper right is visited with nearly vanishing probability.

crease the cardinality (not shown), the largest area state (9) becomes sub-divided based on the ratio $g_{\text{Hb}}/g_{\text{Kr}}$, making it possible to capture the substantial amount of positional information that is available in the region where Hb and Kr expression levels cross, $x/L \sim 0.45$. This information is available only to mechanisms that can measure expression levels with high resolution, and the most efficient use of this resolution is to encode relative levels rather than absolute levels. This need for sensitivity to relative expression levels is the analog for multiple genes of the need for non-monotonicity in the case of a single gene.

In fact, mechanisms that are responsive to combinations of expression levels are quite natural. Enhancer regions typically have closely spaced binding sites for multiple distinct transcription factors [26, 35], and it is easy to imagine that binding events at these different sites interact. Detailed models of the *eve* stripe enhancers show that activation reflects a competition between activation and repression by different gap gene proteins [23, 36], in qualitative agreement with a model in which the state of the enhancers serve as the compressed variable(s) C .

Finally we can ask what happens when we consider all four gap genes together, as in Fig 5. We know that in total there is more than four bits of positional information




FIG. 5: The information bottleneck for positional information carried by all four gap gene expression levels. (A) Expression vs position along the anterior–posterior axis for Hb (red), Kr (blue), Kni (green), and Gt (mustard). Mean (solid) and standard deviation (shading) across $N_{\text{em}} = 38$ embryos in a five minute window (40–44 min) in nuclear cycle 14 [17]. (B) Information bottleneck results, as in Fig 2. Optimal solutions with $\|C\| = 8, 100$ and 800 (shades of grey), and solutions with independent compression of each gene expression level (blue); point corresponds to one bit per gene. (C) Decoding maps $P(x'|x)$ based on compressed representations of the expression levels. No compression (top), $I(C; \{g_i\}) = 4$ bits (middle), and $I(C; \{g_i\}) = 2$ bits (bottom).

available, so that gap gene expression levels specify nuclear positions with $\sim 1\%$ accuracy along the length of the embryo [15]. We can capture a large fraction of this information by keeping only four bits of information about the expression levels, or just one bit per gene (Fig 5B). But even this is not the same as reading expression level individually to one bit precision, since the optimal use of these four bits is to encode combinations of expressions levels, as in the case of two genes in Fig 4.

To extract all the available positional information requires mechanisms that preserve eight or more bits of information about the combined expression levels of the four gap genes. We can visualize what is being gained using the decoding maps introduced in Ref [17]. At the top of fig 5C we consider cells at position x that have access to the expression levels of all four gap genes, and use these data to infer their position x' ; gray levels show the probability distributions $P(x'|x)$, which in this case form a narrow band around $x' = x$, with width $\sigma_{x'}/L \sim 0.01$. In the lower panels we imagine that inference is based not on the actual expression levels but on the compressed version C that emerges from solving the bottleneck problem, and we do this for the optimal compressions with

$I(C; \{g_i\}) = 2$, and 4 bits. We see that as the compression becomes more severe, the inference becomes more uncertain (larger $\sigma_{x'}$) and genuinely ambiguous.

If we demand that compressed variables be constructed from individual gene expression levels, so that $C = \{C_i\}$ and $g_{\text{Hb}} \rightarrow C_1, g_{\text{Kr}} \rightarrow C_2$, etc, then even with same total information capacity we capture less positional information, as indicated by the blue line in Fig 5. This loss of efficiency is substantial, and again indicates the importance of having regulatory mechanisms that are sensitive to combinations of expression levels.

To summarize, individual regulatory mechanisms have limited information capacity [28], and our central result is that this capacity in turn sets strict limits on the amount of positional information that can be extracted from the gap gene expression levels. Extracting all the available information requires that cells have an array of parallel responsive elements. Further, to be efficient each element must respond not to individual gap genes signals but to combinations of these concentrations, so that the space of four concentrations is shattered into compact regions, generalizing the example of two genes shown in Fig 4C. In fact, the readout of positional information encoded in the gap genes is implemented by the array of enhancers controlling pair rule gene expression, and these enhancers are prototypical instances of regulatory elements that respond to combinations of transcription factors [23, 26, 36]. While there is some distance between our abstract formulation and the molecular details, it is attractive to see that this mechanistic complexity is required as a response to basic physical and information theoretic limitations.

We thank P-T Chen, M Levo, R Munshi, and B van Opheusden for helpful discussions. This work was supported in part by the US National Science Foundation, through the Center for the Physics of Biological Function (PHY-1734030) and the Center for the Science of Information (CCF-0939370); by National Institutes of Health Grant R01GM097275; by the Alexander von Humboldt Stiftung; and by the Howard Hughes Medical Institute.

-
- [1] W Bialek and S Setayeshgar, Physical limits to biochemical signaling. *Proc Natl Acad Sci (USA)* **102**, 10040–10045 (2005).
 - [2] G Tkačik, CG Callan Jr, and W Bialek, Information flow and optimization in transcriptional regulation. *Proc Natl Acad Sci (USA)* **105**, 12265–12270 (2008).
 - [3] JO Dubuis, R Samanta and T Gregor, Accurate measurements of dynamics and reproducibility in small genetic networks. *Mol Sys Biol* **9**, 639 (2013).
 - [4] T Gregor, DW Tank, EF Wieschaus, and W Bialek, Probing the limits to positional information. *Cell* **130**, 153–164 (2007).
 - [5] A Abu-Arish, A Porcher, A Czerwonka, N Dostatni and C Fradin, High mobility of Bicoid captured by fluo-

- rescence correlation spectroscopy: Implication for the rapid establishment of its gradient. *Biophys J* **99**, L33–35 (2010).
- [6] JA Drocco, O Grimm, DW Tank and EF Wieschaus, Measurement and perturbation of morphogen lifetime: Effects on gradient shape. *Biophys J* **101**, 1809–1815 (2011).
- [7] CE Hannon, SA Blythe and EF Wieschaus, Concentration dependent chromatin states induced by the bicoid morphogen gradient. *eLife* **6**, e28275 (2017).
- [8] SE Keenan, SA Blythe, RA Marmion, NJ-V Djabrayan, EF Wieschaus, and SY Shvartsman, Rapid dynamics of signal-dependent transcriptional repression by Capicua. *Dev Cell* **52**, 794–801 (2020).
- [9] HC Berg and EM Purcell, Physics of chemoreception. *Biophys J* **20**, 193–219 (1977).
- [10] K Kaizu, WH de Ronde, J Paijmans, K Takahashi, F Tostevin, and PR ten Wolde, The Berg–Purcell limit revisited. *Biophys J* **106**, 976–985 (2014).
- [11] T Friedlander, R Prizak, CC Guet, NH Barton, and G Tkačik, Intrinsic limits to gene regulation by global crosstalk. *Nat Commun* **7**, 12307 (2016).
- [12] T Mora and I Nemenman, Physical limit to concentration sensing in a changing environment. *Phys Rev Lett* **123**, 198101 (2019).
- [13] C Nüsslein–Vollhard and EF Wieschaus, Mutations affecting segment number and polarity in *Drosophila*. *Nature* **287**, 795–801 (1980).
- [14] J Jaeger, The gap gene network. *Cell Mol Life Sci* **68**, 243–274 (2011).
- [15] JO Dubuis, G Tkačik, EF Wieschaus, T Gregor, and W Bialek, Positional information, in bits. *Proc Natl Acad Sci (USA)* **110**, 16301–16308 (2013).
- [16] F Liu, AH Morrison, and T Gregor, Dynamic interpretation of maternal inputs by the *Drosophila* segmentation gene network. *Proc Natl Acad Sci (USA)* **110**, 6724–6729 (2013).
- [17] MD Petkova, G Tkačik, W Bialek, EF Wieschaus, and T Gregor, Optimal decoding of cellular identities in a genetic network. *Cell* **176**, 844–855 (2019).
- [18] L Wolpert, Positional information and the spatial pattern of cellular differentiation. *J Theor Biol* **25**, 1–47 (1969).
- [19] RN Gutenkunst, JJ Waterfall, FP Casey, KS Brown, CR Myers, and JP Sethna, Universally sloppy parameter sensitivities in systems biology models. *PLoS Comput Biol* **3**, e189 (2007).
- [20] MK Transtrum, B Machta, K Brown, BC Daniels, CR Myers, and J Sethna, Sloppiness and emergent theories in physics, biology, and beyond. *J Chem Phys* **143**, 010901 (2015).
- [21] P François and ED Siggia, Predicting embryonic patterning using mutual entropy fitness and in silico evolution. *Development* **137**, 2385–2395 (2010).
- [22] L Bintu, NE Buchler, HG Garcia, U Gerland, T Hwa, J Kondev, and R Phillips, Transcriptional regulation by the numbers: models. *Curr Opin Genet Dev* **15**, 116–124 (2005).
- [23] E Segal, T Raveh-Sadka, M Schroeder, U Unnerstall, and U Gaul, Predicting expression patterns from regulatory sequence in *Drosophila* segmentation. *Nature* **451**, 535–540 (2008).
- [24] G Tkačik, AM Walczak, and W Bialek, Optimizing information flow in small genetic networks. *Phys Rev E* **80**, 031920 (2009); arXiv:0903.4491 [q-bio.MN] (2009).
- [25] O Pulkkinen and R Metzler, Distance matters: The impact of gene proximity in bacterial gene regulation. *Phys Rev Lett* **110**, 198101 (2013).
- [26] EE Furlong and M. Levine, Developmental enhancers and chromosome topology. *Science* **361**, 1341–1345 (2018).
- [27] BR Sabari et al, Coactivator condensation at super-enhancers links phase separation and gene control. *Science* **361** eaar3958 (2018).
- [28] G Tkačik, CG Callan Jr, and W Bialek, Information capacity of genetic regulatory elements. *Phys Rev E* **78**, 011910 (2008).
- [29] N Tishby, FC Pereira, and W Bialek, The information bottleneck method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, B Hajek and RS Sreenivas, eds, pp 368–377 (University of Illinois, 1999); arXiv:physics/0004057 (2000).
- [30] N Slonim and N Tishby, Agglomerative information bottleneck. In *Advances in Neural Information Processing 12*, S Solla, T Leen, and K Müller, eds, pp 617–623 (MIT Press, Cambridge, 2000).
- [31] In molecular biology and genetics, two pathways are redundant if function persists when one is eliminated. In information theory, two pathways are redundant if the information that they carry together is less than the sum of what they carry individually. It is this mathematical notion of redundancy which is relevant here. TM Cover and JA Thomas, *Elements of Information Theory* (Wiley, New York, 1991).
- [32] C Schulz and D Tautz, Autonomous concentration-dependent activation and repression of *Krüppel* by *hunchback* in the *Drosophila* embryo. *Development* **120**, 3043–3049 (1994).
- [33] L Descheemaeker, E Peters, and S de Buyl, Non-monotonic auto-regulation in single gene circuits. *PLoS ONE* **14**, e0216089 (2019).
- [34] MV Staller, BJ Vincent, MDJ Bragdon, T Lydiard-Martin, Z Wunderlich, J Estrada, and AH DePace, Shadow enhancers enable Hunchback bifunctionality in the *Drosophila* embryo. *Proc Natl Acad Sci (USA)* **112**, 785–790 (2015).
- [35] JO Yáñez-Cuna, EZ Kvon, and A Stark, Deciphering the transcriptional cis-regulatory code, *Trends Genet*, **29**, 11–22 (2013).
- [36] J Crocker, GR Ilsley, DL Stern, Quantitatively predictable control of *Drosophila* transcriptional enhancers in vivo with engineered transcription factors *Nat Genet*, **48**, 292–298 (2016).