Deciphering regulatory architectures from synthetic single-cell expression patterns

Rosalind Wenshan Pan[1*], Tom Röschinger[1], Kian Faizi[1], Hernan Garcia[2,3,4,5], Rob Phillips[1,6*]

[1]Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, CA; [2]Biophysics Graduate Group, University of California, Berkeley, CA; [3]Department of Physics, University of California, Berkeley, CA; [4]Department of Molecular and Cell Biology, University of California, Berkeley, CA; [5]Institute for Quantitative Biosciences-QB3, University of California, Berkeley, CA; [6]Department of Physics, California Institute of Technology, Pasadena, CA

* Correspondence: rosalind@caltech.edu and phillips@pboc.caltech.edu

Abstract

For the vast majority of genes in sequenced genomes, there is limited understanding of how they are regulated. Without such knowledge, it is not possible to perform a quantitative theory-experiment dialogue on how such genes give rise to physiological and evolutionary adaptation. One category of high-throughput experiments used to understand the sequence-phenotype relationship of the transcriptome is massively parallel reporter assays (MPRAs). However, to improve the versatility and scalability of MPRA pipelines, we need a "theory of the experiment" to help us better understand the impact of various biological and experimental parameters on the interpretation of experimental data. These parameters include binding site copy number, where a large number of specific binding sites may titrate away transcription factors, as well as the presence of overlapping binding sites, which may affect analysis of the degree of mutual dependence between mutations in the regulatory region and expression levels. To that end, in this paper we create tens of thousands of synthetic single-cell gene expression outputs using both equilibrium and out-of-equilibrium models. These models make it possible to imitate the summary statistics (information footprints and expression shift matrices) used to characterize the output of MPRAs and from this summary statistic to infer the underlying regulatory architecture. Specifically, we use a more refined implementation of the so-called thermodynamic models in which the binding energies of each sequence variant are derived from energy matrices. Our simulations reveal important effects of the parameters on MPRA data and we demonstrate our ability to optimize MPRA experimental designs with the goal of generating thermodynamic models of the transcriptome with base-pair specificity. Further, this approach makes it possible to carefully examine the mapping between mutations in binding sites and their corresponding expression profiles, a tool useful not only for better designing MPRAs, but also for exploring regulatory evolution.

Author summary

With the rapid advancement of sequencing technology, there has been an exponential increase in the amount of data on the genomic sequences of diverse organisms. Nevertheless, deciphering the sequence-phenotype mapping of the genomic data remains a formidable task, especially when dealing with non-coding sequences such as the promoter. In current databases, annotations on transcription factor binding sites are sorely lacking, which creates a challenge for developing a systematic theory of transcriptional regulation. To address this gap in knowledge, high-throughput methods such as massively parallel reporter assays (MPRAs) have been employed to decipher the regulatory genome. In this work, we make use of thermodynamic models to computationally simulate MPRAs in the context of transcriptional regulation and produce thousands of synthetic MPRA datasets. We examine how well typical experimental and data analysis procedures of MPRAs are able to recover common regulatory architectures under different sets of experimental and biological parameters. By establishing a dialogue between high-throughput experiments and a physical theory of transcription, our efforts serve to both improve current experimental procedures and enhancing our broader understanding of the sequence-function landscape of regulatory sequences.

Introduction

With the widespread emergence of sequencing technology, we have seen an explosion of genomic data in recent years. However, data on transcriptional regulation remains far behind. Even for organisms as widely studied as *E. coli*, many promoters lack annotations on the transcription factor binding sites that underlie transcriptional regulation. Moreover, existing binding site annotations are largely without experimental validation for functional activity, as a large proportion are determined through DNA-protein interaction assays such as ChIP-Seq [1, 2, 3, 4] or computational prediction [5]. This fundamental gap in knowledge poses a major obstacle for us to understand the spatial and temporal control of cellular activity, as well as how cells and organisms respond both physiologically and evolutionarily to environmental signals.

One strategy to understand the regulatory genome is by conducting massively parallel reporter assays (MPRAs), where the regulatory activities of a library of sequences are measured simultaneously via a reporter. The library of sequences may be genomic fragments [6] or sequence variants containing mutations relative to the wild-type regulatory sequence [7]. There are two main ways to measure regulatory activities in MPRAs. The first approach uses fluorescence-activated cell sorting to sort cells into bins based on the expression levels of a fluorescent reporter gene [8]. Subsequently, deep sequencing is utilized to determine which sequence variant is sorted into which bin. The second approach uses RNA-sequencing (RNA-Seq) to measure the counts of barcodes associated with each sequence variant as a quantitative read-out for expression levels. The two approaches have been used in both prokaryotic [9, 10, 11, 12] and eukaryotic systems [13, 14, 15] to study diverse genomic elements including promoters and enhancers.

In particular, our group has developed Reg-Seq [16], an RNA-Seq-based MPRA that was used to successfully decipher the regulatory architecture of 100 promoters in *E. coli*, with the hope now to complete the regulatory annotation of entire bacterial genomes. Mutations in regulatory elements lead to reduced transcription factor binding, which may result in measurable changes in expression. Therefore, the key strategy to annotate transcription factor binding sites based on MPRA data that we focus on is to identify sites where mutations have a high impact on expression levels. To do this, one approach is to calculate the mutual information between base identity and expression levels at each site. A so-called information footprint can then be generated by plotting the mutual information at each position along the promoter. Positions with high mutual information are identified as putative transcription factor binding sites.

In this paper, we develop a computational pipeline that simulates the RNA-Seq-based MPRA pipeline. Specifically, we make use of equilibrium statistical mechanics to build synthetic datasets that simulate the experimental MPRA data and examine how various parameters affect the output of MPRAs. Conventionally, thermodynamic models are not sequence specific. The binding energies are usually phenomenological parameters that are fit once and for all. Here, by way of contrast, the idea is to leverage the experimentally

determined or synthetically engineered energy matrices that allow us to consider arbitrary binding site sequences and to compute their corresponding level of expression.

With this computational pipeline, we examine tens of thousands of unique promoters and hence tens of thousands of unique implementations of the sequence-specific thermodynamic models. These sequences are then converted into the two primary summary statistics used to analyze the experimental data, namely, information footprints and expression shift matrices. Given that the experimental implementations of these MPRAs entail tens of millions of unique DNA constructs, the computational pipeline gives us the opportunity to systematically and rigorously analyze the connection between key parameters. These parameters include experimental parameters, such as the rate of mutation used to generate the sequence variants, as well as biological parameters such as transcription factor copy number, where a large number of specific binding sites may titrate away transcription factors. This computational pipeline will help us to optimize MPRA experimental design with the goal of accurately annotating transcription factor binding sites in regulatory elements, while revealing the limits of MPRA experiments in elucidating complex regulatory architectures. Additionally, the insights gained from our simulation platform will enable further dialogue between theory and experiments in the field of transcription including efforts to understand how mutations in the evolutionary context give rise to altered gene expression profiles and resulting organismal fitness.

Our use of thermodynamic models in this pipeline is motivated by several principal considerations. First, because of their simplicity, these models have served and continue to serve as a powerful null model when considering signaling, regulation, and physiology. Their application runs the gamut from the oxygen binding properties of hemoglobin [17, 18, 19, 20, 21], to the functioning of membrane-bound receptors in chemotaxis and quorum sensing [22, 23, 24, 25], and to the binding of transcription factors at their target DNA sequences [26, 27, 28, 29, 30, 31, 32, 33, 34]. Second, for our purposes, the thermodynamic models form an internally consistent closed theoretical system in which we can generate tens of thousands of "single-cell" expression profiles and use the same tools that we use to evaluate real MPRA data to evaluate these synthetic datasets, thus permitting a rigorous means for understanding such real data. That said, despite their many and varied successes, thermodynamic models deserve continued intense scrutiny since many processes of the central dogma involve energy consumption and thus may involve steady state probabilities that while having the form of ratios of polynomials, are no longer of the Boltzmann form [28, 35, 36, 37, 38]. These shortcomings of thermodynamic models are further discussed in S1 Appendix of our Supplemental Information. To address these shortcomings, we also provide a preliminary study of how broken detailed balance might change the interpretation of summary statistics such as information footprints.

The remainder of this paper is organized as follows. In Sec 1.1, we introduce our procedure to construct and analyze synthetic datasets for both promoters regulated by a single transcription factor and promoters regulated by a combination of multiple transcription factors. In Sec 1.2 and Sec 1.3, we discuss the choice of parameters related to the construction of the mutant library, including the rate of mutation, mutational biases, and library size. After setting up the computational pipeline, we perturb biological parameters and examine how these perturbations affect our interpretation of MPRA summary statistics. The parameters that we will explore include the free energy of transcription factor binding (Sec 2.1), the regulatory logic of the promoter (Sec 2.2), the copy number of the transcription factor binding sites (Sec 2.3), and the concentration of the inducers (Sec 2.4). Next, we explore factors that may affect signal-to-noise ratio in the information footprints. These factors include stochastic fluctuations of transcription factor copy number (Sec 3.1), non-specific binding events along the promoter (Sec 3.2), as well as the presence of overlapping binding sites (Sec 3.3). Additionally, in Sec 4, we generalize our pipeline and consider the cases of transcriptional regulation where detailed balance may be broken. Finally, we discuss the insights generated from our computational pipeline in relation to future efforts to decipher regulatory architectures in diverse genomes.

Results

1 Mapping sequence specificity and expression levels

1.1 Computational pipeline for deciphering regulatory architectures from first principles

In MPRA pipelines, the goal is to make the connection between regulatory sequences, transcription factor binding events, and expression levels. In Reg-Seq for example, the authors start with a library of sequence

variants for an unannotated promoter, each of which contains a random set of mutations relative to the wild type sequence. Then, RNA-Seq is used to measure the expression levels of a reporter gene directly downstream of each promoter. By calculating the mutual information between mutations and the measured expression levels, the regulatory architecture of the promoter can be inferred. Finally, Bayesian models and thermodynamic models can be built using statistical mechanics to infer the interaction energies between transcription factors and their binding sites in absolute k_BT units at a base-by-base resolution [16].

Our computational MPRA pipeline involves similar steps, but instead of starting from experimental measurements of expression levels, we use thermodynamic models to predict expression levels given the sequences of the promoter variants and the corresponding interaction energies, as schematized in Fig 1. Through this process, we generate synthetic datasets of expression levels that are in the same format as the datasets that we obtain via RNA-Seq. Subsequently, we analyze the synthetic datasets in the same way as we would analyze an experimental dataset. Importantly, we can perturb various experimental and biological parameters of interest within this pipeline and examine how changing these parameters affect our ability to discover unknown transcription factor binding sites through MPRAs.

We first demonstrate our computational pipeline using a promoter with the simple repression regulatory architecture, i.e. the gene is under the regulation of a single repressor. Specifically, we use the promoter sequence of LacZYA. We assume that it is transcribed by the σ^{70} RNAP and only regulated by the LacI repressor, which binds to the O1 operator within the LacZYA promoter. As we have derived in S2 Appendix, for a gene with the simple repression regulatory architecture, the probability of RNAP being bound [39, 40] is given by

$$p_{\text{bound}} = \frac{\frac{P}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon_{\text{pd}}}}{1 + \frac{P}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon_{\text{pd}}} + \frac{R}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon_{\text{rd}}}}. \tag{1}$$

$\beta = \frac{1}{k_BT}$ where k_B is Boltzmann's constant and T is temperature. As we can see, the parameters that we need are the copy number of RNAP (P), the copy number of repressor (R), the number of non-binding sites (N_{NS}), and the binding energies for RNAP and the repressors ($\Delta\varepsilon_{\text{pd}}$ and $\Delta\varepsilon_{\text{rd}}$), respectively. We begin by assuming that P and R are constant and N_{NS} is the total number of base pairs in the *E. coli* genome. On the other hand, the values for $\Delta\varepsilon_{\text{pd}}$ and $\Delta\varepsilon_{\text{rd}}$ depend on the sequence of the promoter variant.

We calculate the binding energies by mapping the sequences of the promoter variants to the energy matrices of the RNAP and the repressor, as shown in Fig 2(A). Specifically, we assume that binding energies are additive. This means that given a sequence of length l, the total binding energy $\Delta\varepsilon$ can be written as

$$\Delta\varepsilon = \sum_{i=1}^{l} \varepsilon_{i,b_i}, \tag{2}$$

where ε_{i,b_i} is the binding energy corresponding to base identity b_i at position i according to the energy matrix. Here, we use the energy matrices of the RNAP and LacI that were previously experimentally determined using Sort-Seq [41, 42], as shown in Fig 2(B). Unless otherwise specified, these energy matrices are used to build all synthetic datasets in the remainder of this paper. It should be acknowledged that the additive model does not take into account epistasis effects [43, 44]. It may be beneficial to include higher-order interaction energy terms in future simulations of MPRA pipelines.

After computing the sequence-specific binding energies, we can then substitute the relevant energy terms into Eq 1 and calculate the probability of RNAP being bound. To connect the probability of RNAP being bound to expression levels, we make use of the occupancy hypothesis, which states that the rate of mRNA production is proportional to the probability of RNAP occupancy at the promoter [45]. The rate of change in mRNA copy number is given by the difference between the rates of mRNA production and degradation. In general, there can be multiple transcriptionally active states, each with its own transcription rate. For example, for a promoter that is regulated by an activator, there are two transcriptionally active states. One is where only RNAP is bound to the promoter and one is where both RNAP and the activator are bound to the promoter. Each of these two states could have a different rate of mRNA production. With this, we define an average rate of mRNA production, which is given by the sum of each state's production rate, weighted by

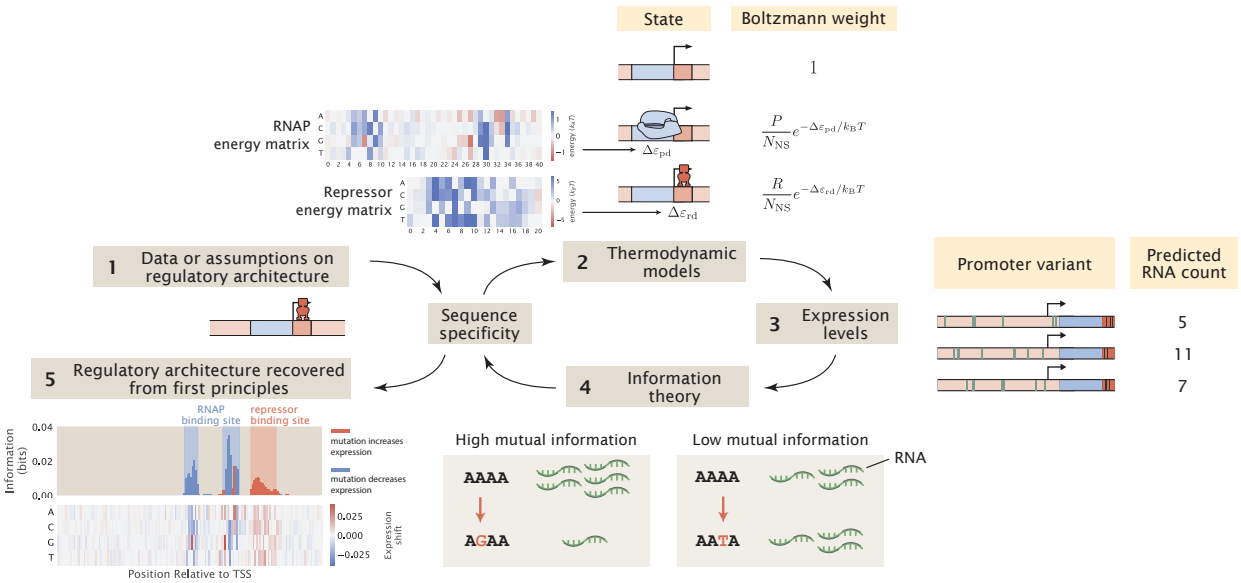

Fig 1. A computational pipeline for deciphering regulatory architectures from first principles. Given (1) knowledge or assumptions about the regulatory architecture of a promoter, we make use of (2) thermodynamic models to construct a states-and-weights diagram, which contains information about all possible states of binding and the associated Boltzmann weights. Here, in the states-and-weights diagram, P is the copy number of RNAP, R is the copy number of the repressor, N_{NS} is the number of non-binding sites, $\Delta\varepsilon_{\mathrm{pd}}$ and $\Delta\varepsilon_{\mathrm{rd}}$ represent the binding energies of RNAP and the repressors at their specific binding sites relative to the non-specific background, respectively. Using these states-and-weights diagrams as well as the energy matrices, which are normalized to show the change in binding energies for any mutation along the promoter compared to the wild-type sequence, we can (3) predict the expression levels for each of the promoter variants in a mutant library. To recover the regulatory architecture, we (4) calculate the mutual information between the predicted expression levels and mutations at each position along the promoter according to Eq 6. In particular, there is high mutual information if a mutation leads to a large change in expression and there is low mutual information if a mutation does not lead to a significant change in expression. The mutual information at each position is plotted in an information footprint, where the height of the peaks corresponds to the magnitude of mutual information, and the peaks are colored based on the sign of expression shift, defined in Eq 9. Given the assumption that the positions with high mutual information are likely to be RNAP and transcription factor binding sites, we (5) recover the regulatory architecture of the promoter. The base-specific effects of mutations on expression levels can also be seen from expression shift matrices, which is calculated using Eq 10, where the difference between the expression levels of sequences carrying a specific mutation at a given position and the average expression level across all mutant sequences is computed.

the probability of the state $\langle r \rangle = \sum_i r_i p_{\mathrm{bound},i}$. Hence, the rate of change of mRNA copy number is given by

$$\frac{\mathrm{d}m}{\mathrm{d}t} = \sum_i r_i p_{\mathrm{bound},i} - \gamma m, \tag{3}$$

where for transcriptionally active state i, r_i is the rate of transcription, $p_{\mathrm{bound},i}$ is the probability of RNAP occupancy in state i, m is the copy number of mRNAs, and γ is the rate of mRNA degradation. Therefore, the steady-state level of mRNA is given by

$$m^* = \frac{1}{\gamma} \sum_i r_i p_{\mathrm{bound},i}. \tag{4}$$

For simplicity, we assume that each transcriptionally active state has the same rate of mRNA production, r.

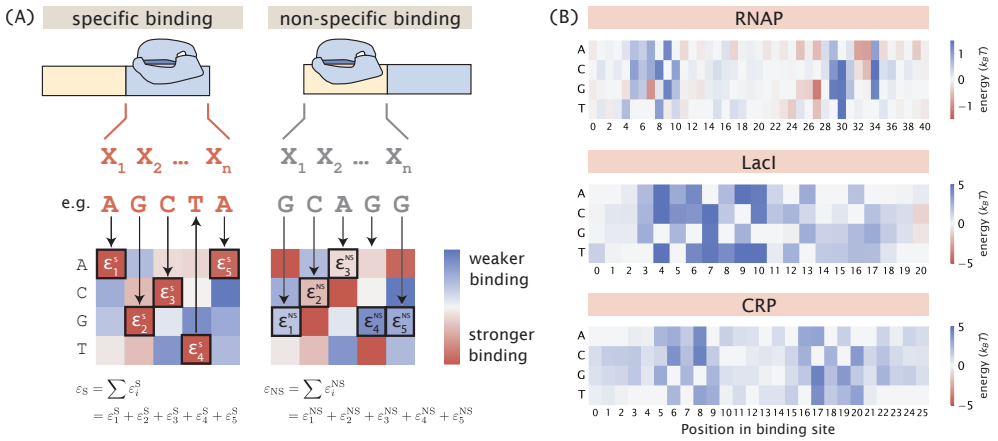

Fig 2. Mapping binding site sequences to binding energies using energy matrices. (A) Given the assumption that binding energies are additive, we can use an energy matrix to determine how much energy each base along the binding site contributes and compute the total binding energy by taking the sum of the binding energies contributed by each position. The total binding energy can be used to compute the Boltzmann weight for each of the states, which is then used to calculate the probability of RNAP being bound. (B) Experimentally measured energy matrices of RNAP [41], the LacI repressor [42], and the CRP activator. [8]

Therefore,

$$m^* = \alpha \sum_i p_{\text{bound},i}, \text{ where } \alpha = \frac{r}{\gamma}. \tag{5}$$

Using the above expression, we can calculate the expected RNA count for each of the promoter variants in our library. Assuming that r and γ do not depend on the mRNA sequence, the total probability of RNAP being bound, given by $p_{\text{bound}} = \sum p_{\text{bound},i}$, is scaled by the same constant to produce the mRNA count of each promoter variant. Therefore, the choice of α does not affect our downstream calculations involving the probability distribution of expression levels. Depending on the sequencing depth, the mRNA count for each promoter variant is typically on the order of 10 to 10^3 [16]. Here, we take the geometric mean and set α to 10^2 to ensure that the mRNA count is on a realistic scale.

Up until this point, we have constructed a synthetic RNA-Seq dataset containing the predicted expression levels of each sequence variant in a mutant library. MPRA data is often described using several interesting summary statistics. Using thousands of synthetically derived mRNA counts, we can compute such summary statistics and ask how both biological and experimental parameters change them. These summary statistics can then be used to infer the underlying regulatory architecture from these large-scale synthetic datasets. To do this, we calculate the mutual information between mutations and expression levels. The mutual information at position i is given by

$$I_i = \sum_b \sum_\mu \Pr{}_i(b,\mu) \log_2 \left(\frac{\Pr{}_i(b,\mu)}{\Pr{}_i(b)\Pr(\mu)} \right), \tag{6}$$

where b represents base identity, μ represents expression level, $\Pr{}_i(b)$ is the marginal probability distribution of mutations at position i, $\Pr(\mu)$ is the marginal probability distribution of expression levels across all promoter variants, and $\Pr{}_i(b,\mu)$ is the joint probability distribution between expression levels and mutations at position i. In general, b can be any of the four nucleotides, i.e. $b \in \{\text{A}, \text{C}, \text{G}, \text{T}\}$. This means that $\Pr{}_i(b)$ is obtained by computing the frequency of each base per position. Alternatively, a more coarse grained approach can be taken, where the only distinction is between the wild-type base and mutation, in which case b is defined as

$$b = \begin{cases} 0, & \text{if the base is mutated,} \\ 1, & \text{if the base is wild type.} \end{cases} \tag{7}$$

As shown in S4 Appendix, using the coarse-grained definition of b improves the signal-to-noise ratio of the information footprint as reducing the number of states reduces articifical noise outside of the specific binding sites. Therefore, we use this definition for our subsequent analysis.

On the other hand, to represent expression levels as a probability distribution, we group sequences in each range of expression levels into discrete bins and compute the probabilities that a given promoter variant is found in each bin. As shown in S4 Appendix, we found that increasing the number of bins leads to a lower signal-to-noise ratio in the information footprints because the additional bins contribute to artificial noise. Therefore, we choose to use only two bins with the mean expression level as the threshold between them. This means that μ can take the values of

$$\mu = \begin{cases} 0, & \text{if expression is lower than mean expression} \\ 1, & \text{if expression is higher than mean expression.} \end{cases} \tag{8}$$

In S5 Appendix, we derive the information footprint for a constitutive promoter analytically and demonstrate that in the absence of noise, mutual information is expected to be 0 outside of the specific binding sites and non-zero at a specific binding site.

To decipher the regulatory architecture of a promoter, another important piece of information is the direction in which a mutation changes expression. This can be determined by calculating the expression shift, which measures the change in expression when there is a mutation at a given position [46]. Suppose there are n promoter variants in our library, then the expression shift Δs_l at position l is given by

$$\Delta s_l = \frac{1}{n} \sum_{i=1}^{n} \xi_{i,l} \left(c_i - \langle c \rangle \right), \text{ where } \langle c \rangle = \frac{1}{n} \sum_{i=1}^{n} c_i, \tag{9}$$

where c_i represents the RNA count of the i-th promoter variant, $\xi_{i,l} = 0$ if the base at position l in the i-th promoter variant is wild type, and $\xi_{i,l} = 1$ if the base is mutated. If the expression shift is positive, it indicates that mutations lead to an increase in expression and the site is likely to be bound by a repressor. On the other hand, a negative expression shift indicates that mutations lead to a decrease in expression, and therefore the site is likely to be bound by RNAP or an activator.

By calculating the mutual information and expression shift at each base position along the promoter, we can plot an information footprint for a promoter with the simple repression regulatory architecture, as shown in Fig 3(B). There are two peaks with negative expression shifts near the -10 and -35 positions, which correspond to the canonical RNAP binding sites. There is another peak immediately downstream from the transcription start site with a positive expression shift, which corresponds to the binding site of the LacI repressor. Taken together, we have demonstrated that by calculating mutual information, we are able to recover binding sites from our synthetic dataset on expression levels.

To have a more precise understanding on how much each mutation to each possible base identity $b \in \mathrm{A, C, G, T}$ changes expression levels, we can extend Eq 9 to calculate an expression shift matrix. Specifically, the value in the expression shift matrix at position l corresponding to base b is given by

$$\Delta s_{b,l} = \begin{cases} \frac{1}{n} \sum_{i=1}^{n} \xi_{b,i,l} \left(\frac{c_i}{\langle c \rangle} - 1 \right), \text{ where } \langle c \rangle = \frac{1}{n} \sum_{i=1}^{n} c_i, & \text{if } b \text{ is mutated} \\ 0, & \text{if } b \text{ is wild type} \end{cases} \tag{10}$$

where $\xi_{b,i,l} = 1$ if the base at position l in the i-th promoter variant corresponds to base identity b and $\xi_{b,i,l} = 0$ otherwise. Note that in comparison to Eq 9, here we are calculating the relative change in expression, which is easier to interpret than the absolute change in expression. An example of an expression shift matrix is shown in panel (5) of Fig 1.

Using the same procedure as described above, we can also produce synthetic datasets for other classes of regulatory architectures. In Fig 3, we demonstrate that we can recover the expected binding sites based on synthetic datasets for six common types of regulatory architectures [16]. The states-and-weights diagrams and p_{bound} expressions used to produce these synthetic datasets are shown in S3 Appendix.

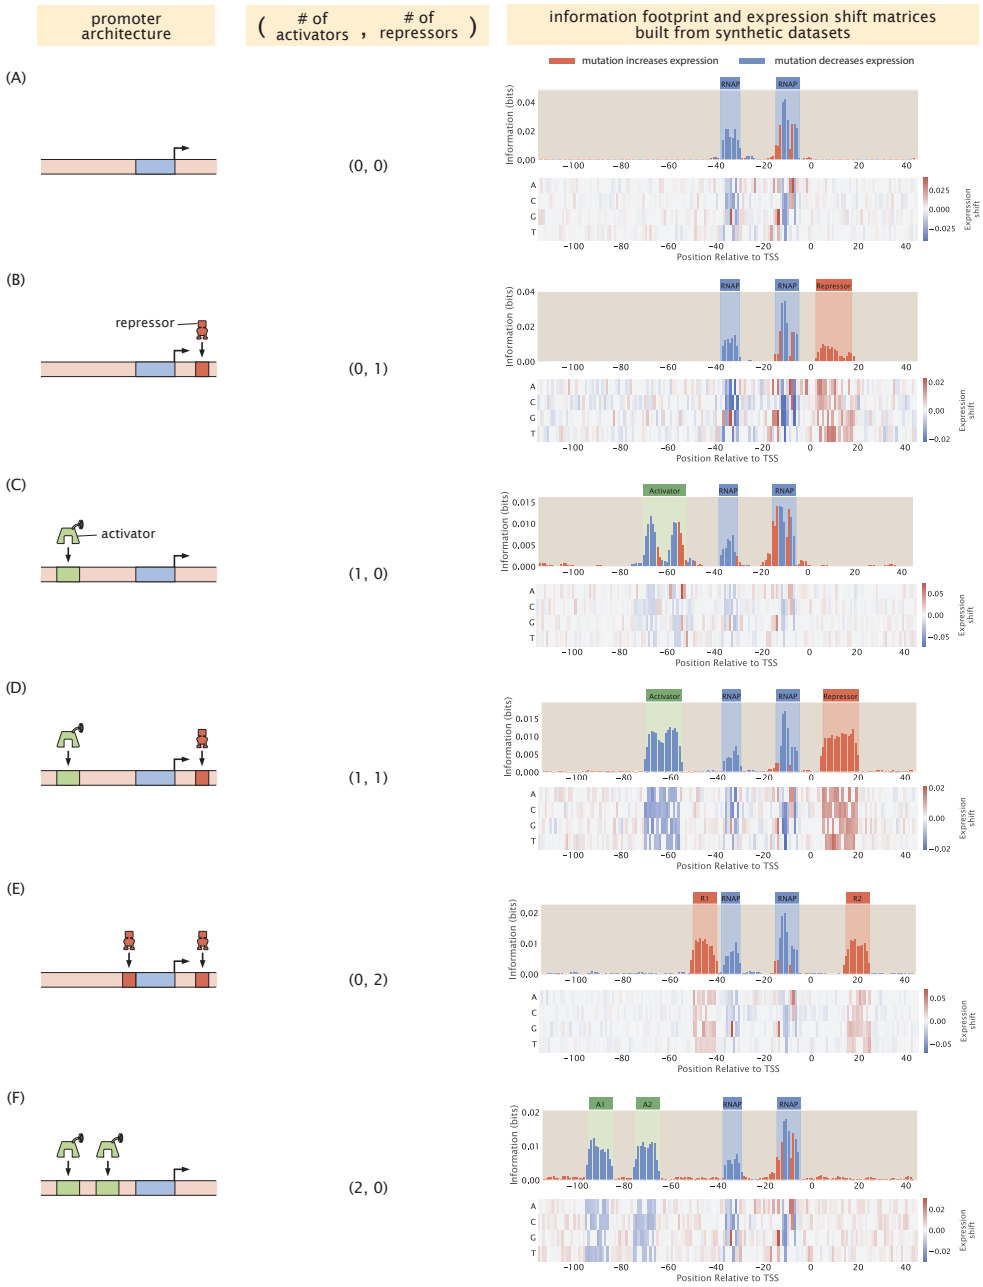

Fig 3. Building information footprints and expression shift matrices based on synthetic datasets of different regulatory architectures. We describe each of the regulatory architectures using the notation (A,R), where A refers to the number of activator binding sites and R refers to the number of repressor binding sites. The corresponding information footprints and expression shift matrices built from synthetic datasets are shown on the right. The architectures shown in panels (A)-(F) are a constitutive promoter, simple repression, simple activation, repression-activation, double repression, and double activation, respectively. For panels (A)-(C), we use energy matrices of RNAP, LacI, and CRP shown in Fig 2(B). For panels (D)-(F), we continue to use the experimentally measured energy matrix for RNAP; the energy matrices for the repressors and the activators are constructed by hand, where the interaction energies at the wild type bases are set to 0 k_BT and the interaction energies at the mutant bases are set to 1 k_BT.

1.2 Changing mutation rates and adding mutational biases

One key parameter in the MPRA pipeline is the level of mutation for each sequence variant in the library. Here, we again consider a gene with the simple repression regulatory architecture as a case study and we examine how varying mutation rates and mutational biases changes the signals in the information footprints. We quantify the level of signal, S, by calculating the average mutual information at each of the binding sites. This is given by

$$S = \langle I \rangle_B = \frac{1}{l} \sum_{i \in B} I_i, \tag{11}$$

where B represents the set of bases within a given binding site, I_i represents the mutual information at base position i, and l is the length of B, i.e. the number of bases in the binding site.

As shown in Fig 4(A) and 4(B), in general, when there is a higher rate of mutation, the average mutual information at the RNAP binding site increases relative to the average mutual information at the repressor binding site. To explain this effect, we consider κ, the ratio between the Boltzmann weights of the repressor and RNAP

$$\kappa = \frac{R \cdot e^{-\beta(\Delta\varepsilon_{\mathrm{rd}} + m_r \Delta\Delta\varepsilon_{\mathrm{rd}})}}{P \cdot e^{-\beta(\Delta\varepsilon_{\mathrm{pd}} + m_p \Delta\Delta\varepsilon_{\mathrm{pd}})}} = \frac{R}{P} \cdot e^{-\beta(\Delta\varepsilon_{\mathrm{rd}} - \Delta\varepsilon_{\mathrm{pd}} + m_r \Delta\Delta\varepsilon_{\mathrm{rd}} - m_p \Delta\Delta\varepsilon_{\mathrm{pd}})}, \tag{12}$$

where m_r and m_p are the number of mutations at the repressor and RNAP binding sites, and $\Delta\Delta\varepsilon_{\mathrm{rd}}$ and $\Delta\Delta\varepsilon_{\mathrm{pd}}$ are the change in binding energies due to each mutation at the repressor and RNAP binding sites. To express κ as a function of the mutation rate θ, we can rewrite m_r and m_p as a product of θ and the lengths of repressor and RNAP binding sites l_r and l_p,

$$\kappa = \frac{R}{P} \cdot e^{-\beta E}, \text{ where } E = \Delta\varepsilon_{\mathrm{rd}} - \Delta\varepsilon_{\mathrm{pd}} + \theta(l_r \Delta\Delta\varepsilon_{\mathrm{rd}} - l_p \Delta\Delta\varepsilon_{\mathrm{pd}}) \tag{13}$$

We assume that $\Delta\Delta\varepsilon_{\mathrm{rd}}$ and $\Delta\Delta\varepsilon_{\mathrm{pd}}$ are equal to the average effect of mutations per base pair within each binding site, which can be calculated using the formula

$$\Delta\Delta\varepsilon = \frac{1}{3l} \sum_{i=1}^{l} \sum_{b \neq b_i}^{\Lambda} \varepsilon_{i,b}, \text{ where } \Lambda = \{\mathrm{A}, \mathrm{T}, \mathrm{C}, \mathrm{G}\}, \tag{14}$$

where $\varepsilon_{i,b}$ is the energy contribution from position i when the base identity is b. As we are using energy matrices where the energies corresponding to the wild-type base identities, b_i, are set to 0, we only need to compute the sum of the energy terms for the mutant bases $b \neq b_i \in \Lambda$ at each position. Since there are three possible mutant bases at each site, it follows that to find the average effect of mutations, we divide the sum of the energy matrix by 3 times the length of the binding site l.

By applying this formula to the energy matrices in Fig 2(B), we see that $\Delta\Delta\varepsilon_{\mathrm{rd}} \approx 2.24\ k_B T$ and $\Delta\Delta\varepsilon_{\mathrm{pd}} \approx 0.36\ k_B T$, where $\Delta\Delta\varepsilon_{\mathrm{pd}}$ is averaged over the 20 bases surrounding the -35 and -10 binding sites. Moreover, $l_r = l_p \approx 20$ base pairs. Therefore,

$$l_r \Delta\Delta\varepsilon_{\mathrm{rd}} - l_p \Delta\Delta\varepsilon_{\mathrm{pd}} = 20 \times 2.24\ k_B T - 20 \times 0.36\ k_B T = 37.6\ k_B T. \tag{15}$$

Since the above value is positive, κ decreases with increasing mutation rate, making the repressor bound state less likely compared to the RNAP bound state. With the repressor bound state becoming less likely, the signal in the repressor binding site goes down, since mutations changing the binding energy of the repressor change the transcription rate less significantly. As shown in S6 Appendix, when we reduce the effect of mutations on the binding energy of the repressor, we recover the signal at the repressor binding site. Conversely, the average mutual information at the repressor binding site increases when the rate of mutation is decreased. This is because when there are very few mutations, the energy E will be less than 0 and therefore κ will be greater than 1. As a result, the repressor will be preferentially bound, which blocks RNAP binding and leads to a low signal at the RNAP binding site. We can recover the signal at the RNAP binding site by increasing the binding energy between RNAP and the wild type promoter, as shown in S6 Appendix.

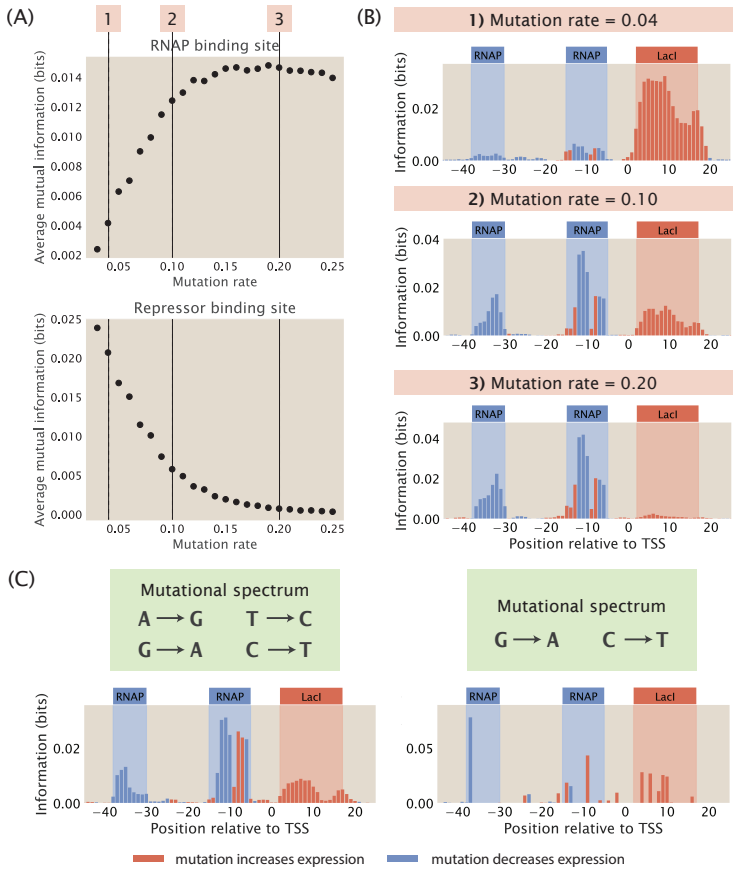

Fig 4. Changing mutation rate and adding mutational biases. (A) Changes in the average mutual information at the RNAP and at the repressor binding sites when the mutation rate of the mutant library is increased. Average mutual information is calculated according to Eq 11. Each data point is the mean of average mutual information across 20 synthetic datasets with the corresponding mutation rate. The numbered labels correspond to information footprints shown in (B). (B) Representative information footprints built from synthetic datasets with mutation rates of 0.04, 0.1, and 0.2. (C) Information footprints built from synthetic datasets where the mutant library has a limited mutational spectrum. The left panel shows a footprint where mutations from A to G, G to A, T to C, and C to T are allowed. The right panel shows a footprint where only mutations from G to A and from C to T are allowed.

Importantly, the use of κ lays the groundwork for finding the optimal rate of mutation in a mutant library. Specifically, we would like to determine the choice of mutation rate that will give us high and balanced signals for both the RNAP and transcription factor binding sites in the information footprints. To find the optimal rate of mutation, we need to satisfy the condition

$$\kappa = \frac{R}{P} e^{-\beta(\Delta\varepsilon_{\mathrm{rd}} - \Delta\varepsilon_{\mathrm{pd}} + \theta(l_r \Delta\Delta\varepsilon_{\mathrm{rd}} - l_p \Delta\Delta\varepsilon_{\mathrm{pd}}))} = 1, \tag{16}$$

which puts repressor and RNAP binding on an equal footing. Plugging in the values of R, P, and the energy terms and solving for θ, we get that $\theta \approx 0.10$. This shows that an intermediate mutation rate is optimal for maintaining high signals at all binding sites. It should be acknowledged that the optimal mutation rate is sensitive to the underlying parameters such as transcription factor copy numbers and binding energies. In S7 Appendix, we numerically compute the optimal mutation rate for parameters within physiological ranges, and we found that the optimal mutation rate can vary from 0.01 to 0.50. Nevertheless, for the remaining analysis shown in this work, we fix the mutation rate at 10%, which is similar to the mutation rate typically used in MPRAs such as Sort-Seq [8] and Reg-Seq [16].

In addition to mutation rate, another important variation in the design of the mutant library is the presence of mutational biases. For example, some mutagenesis techniques, including CRISPR-Cas9, often carry mutational biases whereby mutations within the family of purines and the family of pyrimidines have a higher efficiency compared to mutations between purines and pyrimidines [47]. We build mutant libraries that incorporate two different mutational spectrums. In the first case, we allow only swaps between A and G and between C and T. For this library, we observe that the signals at both the RNAP binding site and the repressor binding site are well preserved, as shown in the left panel of Fig 4(C). In the second case, we only allow mutations from G to A and from C to T without allowing the reverse mutations. As shown in the right panel of Fig 4(C), due to only two bases being allowed to mutate, only a few, possibly low-effect mutations are observed, making small regions such as the -10 and -35 sites hard to detect. These results show that information footprints are robust to mutational biases provided that most sites are allowed to mutate.

1.3 Noise as a function of library size

Another parameter that is important for library design is the total number of sequence variants in the mutant library. We build synthetic datasets with varying library sizes and computed the information footprints. To quantify the quality of signal in information footprints, we calculate signal-to-noise ratio, σ, according to the formula

$$\sigma = \frac{\langle I \rangle_B}{\langle I \rangle_{NB}}, \text{ where } \langle I \rangle_B = \frac{1}{l_B} \sum_{i \in B} I_i \text{ and } \langle I \rangle_{NB} = \frac{1}{l_{NB}} \sum_{i \in NB} I_i. \tag{17}$$

Here, I_i represents the mutual information at position i, B is the set of bases within each binding site, NB is the set of bases outside the binding sites, l_B is the length of the specific binding site, and l_{NB} is the total length of the non-binding sites. As shown in Fig 5(A) and 5(B), we observe that signal-to-noise ratio increases as the library size increases. This may be explained by the "hitch-hiking" effect: since mutations are random, mutations outside of specific binding sites can co-occur with mutations in specific binding sites. As a result, when the library is small, there is an increased likelihood that a mutation outside of specific binding sites and a mutation at a specific binding site become correlated by chance, leading to artificial signal at the non-binding sites.

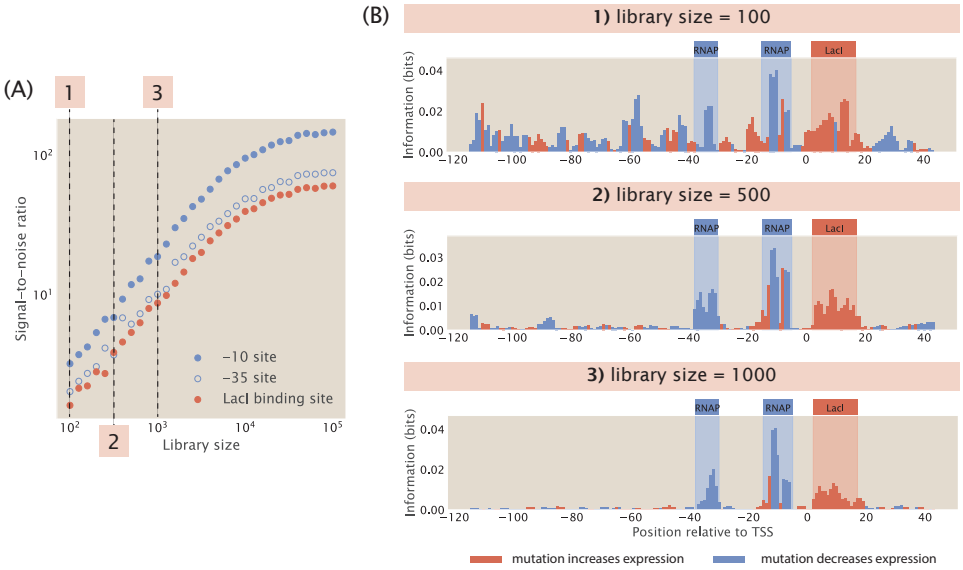

Fig 5. Noise as a function of library size. (A) Signal-to-noise ratio increases as library size increases. Signal-to-noise ratio is calculated according to Eq 17. Each data point is the mean of average mutual information across 20 synthetic datasets with the corresponding library size. The numbered labels correspond to footprints in (B). (B) Representative information footprints with a library size of 100, 500, and 1000.

To demonstrate the hitch-hiking effect analytically, we consider a hypothetical promoter that is constitutively transcribed and only two base pairs long, as illustrated in Fig 6(A). Without loss of generality, we assume that there are only two letters in the nucleotide alphabet, X and Y. Therefore, a complete and unbiased library contains four sequences: XX, YX, XY, and YY. We designate that $\varepsilon_X < \varepsilon_Y$, i.e. the RNAP is strongly bound at the binding site when the base identity is X and weakly bound when the base identity is Y. We also assume that there is active transcription only when RNAP is bound to the second site. Under these assumptions, there are high expression levels when the promoter sequence is XX or YX and low expression levels when the promoter sequence is XY and YY.

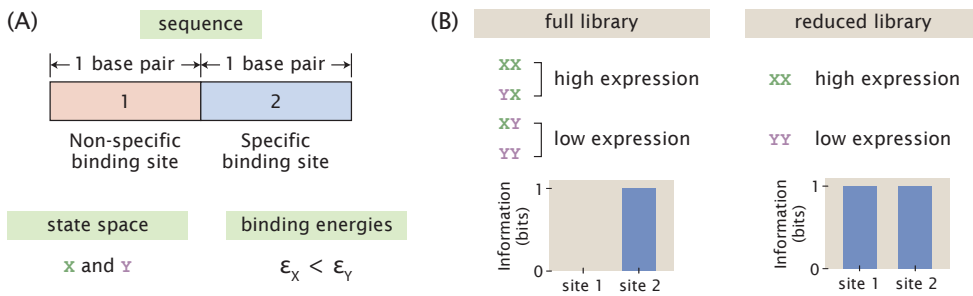

Fig 6. Hitch-hiking effect in the hypothetical minimal promoter. (A) Set-up of the hypothetical minimal promoter. The specific binding sites and the non-binding sites of the minimal promoter are each 1 base-pair long. There are two possible bases at each binding site, X and Y. Strong binding occurs when the base is X, whereas weak binding occurs when the base is Y. (B) Effect of library size on the information footprint of the minimal promoter. A full mutant library consists of all four possible sequences and leads to a footprint with no signal outside of the specific binding sites. On the other hand, a reduced mutant library with only two sequences creates noise outside of the specific binding sites. In this case, the noise at the non-binding site has the same magnitude as the signal at the specific binding site.

We first consider a mutant library with full diversity and no bias, i.e. the four possible sequences, XX, YX, XY, and YY, are each present in the library exactly once. The marginal probability distribution for expression levels is

$$\Pr(\mu) = \begin{cases} 0.5, & \text{if } \mu = 0 \\ 0.5, & \text{if } \mu = 1. \end{cases} \tag{18}$$

The marginal probability distributions of base identity at the two sites are

$$\Pr{}_1(b) = \Pr{}_2(b) = \begin{cases} 0.5, & \text{if } b = \text{X} \\ 0.5, & \text{if } b = \text{Y}. \end{cases} \tag{19}$$

The joint probability distribution at the first site is

$$\Pr{}_1(\mu, b) = \begin{cases} 0.25, & \text{if } \mu = 0 \text{ and } b = \text{X} \\ 0.25, & \text{if } \mu = 0 \text{ and } b = \text{Y} \\ 0.25, & \text{if } \mu = 1 \text{ and } b = \text{X} \\ 0.25, & \text{if } \mu = 1 \text{ and } b = \text{Y}. \end{cases} \tag{20}$$

On the other hand, the joint probability distribution at the second site is

$$\Pr{}_2(\mu, b) = \begin{cases} 0, & \text{if } \mu = 0 \text{ and } b = \text{X} \\ 0.5, & \text{if } \mu = 0 \text{ and } b = \text{Y} \\ 0.5, & \text{if } \mu = 1 \text{ and } b = \text{X} \\ 0, & \text{if } \mu = 1 \text{ and } b = \text{Y}. \end{cases} \tag{21}$$

We can calculate the mutual information at each site according to Eq 6,

$$I_1 = 4\left(\frac{1}{4}\log_2\left(\frac{1/4}{1/2\cdot 1/2}\right)\right) = 0 \tag{22}$$

$$I_2 = 2\left(\frac{1}{2}\log_2\left(\frac{1/2}{1/2\cdot 1/2}\right)\right) = 1. \tag{23}$$

Therefore, when the library has the maximum size, there is perfect signal at the specific binding site and no signal outside of the specific binding site, as shown in Fig 6(B).

On the other hand, consider a reduced library that only consists of XX and YY. According to the assumptions stated above, XX has high expression and YY has low expression. In this case, there is an apparent correlation between the base identity at the non-binding site and expression levels, where a base identity of X at the non-binding site appears to lead to high expression levels and a base identity of Y at the non-binding site appears to lead to low expression levels. To demonstrate this analytically, we again write down the relevant probability distributions required for calculating mutual information. The marginal probability distributions for expression levels and base identity are the same as the case where we have a full library. However, the joint probability distributions at both of the two sites become

$$\Pr{}_1(\mu, b) = \Pr{}_2(\mu, b) = \begin{cases} 0, & \text{if } \mu = 0 \text{ and } b = \text{X} \\ 0.5, & \text{if } \mu = 0 \text{ and } b = \text{Y} \\ 0.5, & \text{if } \mu = 1 \text{ and } b = \text{X} \\ 0, & \text{if } \mu = 1 \text{ and } b = \text{Y}. \end{cases} \tag{24}$$

This means that for both the non-binding site and the specific binding site, the mutual information is

$$I_1 = I_2 = 2\left(\frac{1}{2}\log_2\left(\frac{1/2}{1/2\cdot 1/2}\right)\right) = 1. \tag{25}$$

As shown in Fig 6(B), this creates an artificial signal, or noise, outside of the specific binding sites that cannot be distinguished from the signal at the specific binding site.

For the remaining analyzes shown in this work, we use a library size of 5,000 in order to minimize noise from hitch-hiking effects. We choose not to use a larger library because it would significantly increase the computational cost during parameter searches. Moreover, we would like to use a library size that is experimentally feasible. In Reg-Seq, the average library size is 1,500 [16]. A larger library would make MPRAs cost prohibitive as a high-throughput method.

2 Perturbing biological parameters in the computational pipeline

2.1 Tuning the free energy of transcription factor binding

So far, we have demonstrated that we can build synthetic datasets for the most common regulatory architectures and we have chosen the appropriate mutation rate and library size to construct mutant libraries. Next, we proceed to perturb parameters that affect the probability of RNAP being bound and observe the effects of these perturbations. These analyzes will elucidate the physiological conditions required for obtaining clear signals from transcription factor binding events and delineate the limits of MPRA procedures in identifying unannotated transcription factor binding sites.

We again begin by considering the promoter with the simple repression motif, for which the probability of RNAP being bound is given by Eq 1. It is known that in *E. coli* grown in minimal media, the copy number of RNAP is $P \approx 10^3$ [48, 49] and the copy number of the repressor is $R \approx 10$ [50]. The binding energy of RNAP is $\Delta\varepsilon_{\text{pd}} \approx -5\ k_BT$ [41] and the binding energy of the repressor is $\Delta\varepsilon_{\text{rd}} \approx -15\ k_BT$ [51]. Moreover, assuming that the number of non-binding sites is equal to the size of the *E. coli* genome, we have that $N_{\text{NS}} \approx 4 \times 10^6$. Given these values, we can estimate that

$$\frac{P}{N_{\text{NS}}} e^{-\beta\Delta\varepsilon_{\text{pd}}} \approx \frac{10^3}{4\times 10^6} \times e^5 \approx 0.04 \tag{26}$$

and

$$\frac{R}{N_{\rm NS}} e^{-\beta \Delta\varepsilon_{\rm rd}} \approx \frac{10}{4 \times 10^6} \times e^{15} \approx 8. \tag{27}$$

Since $\frac{P}{N_{\rm NS}} e^{-\beta \Delta\varepsilon_{\rm pd}} \ll 1$, we can neglect this term from the denominator in Eq 1 and simplify $p_{\rm bound}$ for the simple repression motif to

$$p_{\rm bound} = \frac{\frac{P}{N_{\rm NS}} e^{-\beta \Delta\varepsilon_{\rm pd}}}{1 + \frac{R}{N_{\rm NS}} e^{-\beta \Delta\varepsilon_{\rm rd}}}. \tag{28}$$

Furthermore, we define the free energy of RNAP binding as

$$F_P = \Delta\varepsilon_{\rm pd} - k_B T \ln \frac{P}{N_{\rm NS}} \tag{29}$$

and the free energy of repressor binding as

$$F_R = \Delta\varepsilon_{\rm rd} - k_B T \ln \frac{R}{N_{\rm NS}}. \tag{30}$$

Both expressions are written according to the definition of Gibbs free energy, where the first terms correspond to enthalpy and the second terms correspond to entropy. Using these definitions, we can rewrite $p_{\rm bound}$ as

$$p_{\rm bound} = \frac{e^{-\beta F_P}}{1 + e^{-\beta F_R}}. \tag{31}$$

In this section, we specifically examine the changes in the information footprints when we tune F_R. As shown in Fig 7(A) and 7(C), if we increase F_R by reducing the magnitude of $\Delta\varepsilon_{\rm rd}$ or reducing the copy number of the repressor, we lose the signal at the repressor binding site. For example, compared to the O1 operator, the LacI repressor has weak binding energy at the O3 operator, where $\Delta\varepsilon_{\rm rd} \approx -10\ k_B T$ [51]. Therefore

$$F_R = \Delta\varepsilon_{\rm rd} - k_B T \ln \frac{R}{N_{\rm NS}} \approx (-10 - \ln \frac{10}{4 \times 10^6})\ k_B T \approx 3\ k_B T, \tag{32}$$

and

$$e^{-\beta F_R} \approx 0.05. \tag{33}$$

In these cases, $e^{-\beta F_R} \ll 1$ and therefore $e^{-\beta F_R}$ can be neglected from the denominator and the probability of RNAP being bound can be simplified to

$$p_{\rm bound} = e^{-\beta F_P}. \tag{34}$$

This implies that mutations at the repressor binding sites will not have a large effect on $p_{\rm bound}$ and the mutual information at the repressor binding site will be minimal.

On the other hand, if we decrease the free energy of binding either by increasing the magnitude of $\Delta\varepsilon_{\rm rd}$ or increasing the copy number of the repressor, it leads to a stronger signal at the repressor binding site while significantly reducing the signal at the RNAP binding site, as we can see in Fig 7(A) and 7(C). For example, when $\Delta\varepsilon_{\rm rd} = -20 k_B T$, we have that

$$F_R = \Delta\varepsilon_{\rm rd} - k_B T \ln \frac{R}{N_{\rm NS}} \approx (-20 - \ln \frac{10}{4 \times 10^6})\ k_B T \approx -7\ k_B T, \tag{35}$$

and therefore

$$e^{-\beta F_R} \approx 10^3. \tag{36}$$

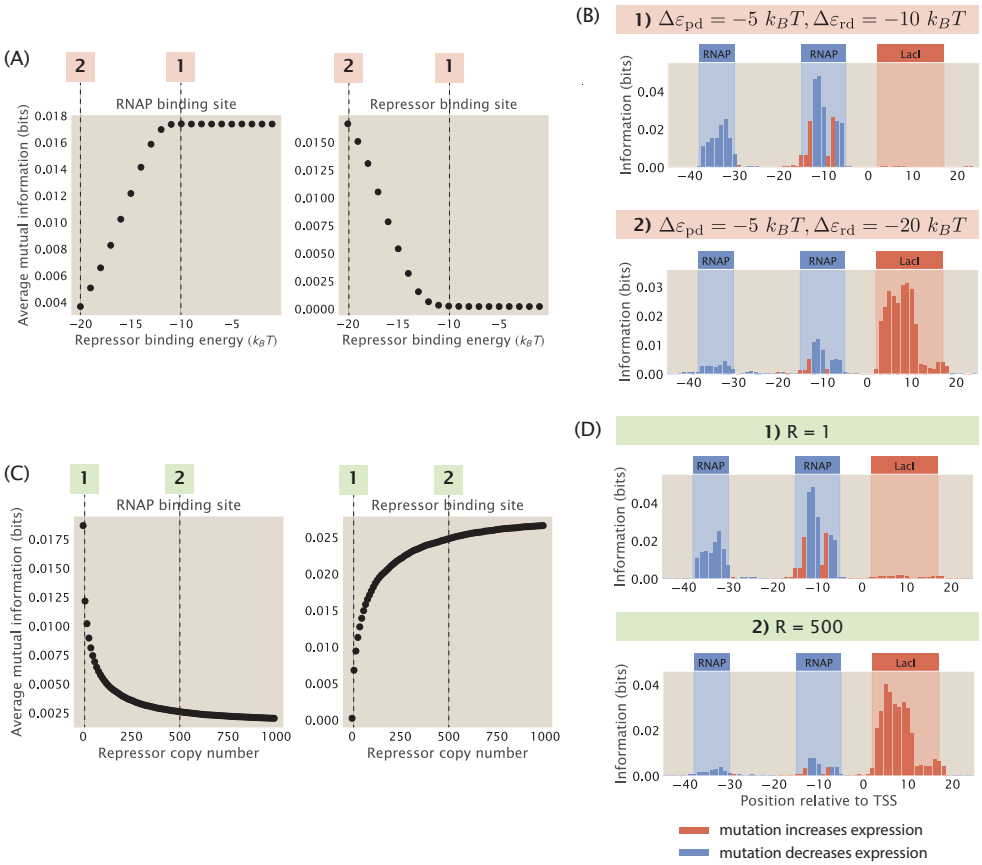

Fig 7. The strength of the signal at binding sites depends on the free energy of repressor binding. (A) Increasing the binding energy of the repressor leads to an increase in average mutual information at the RNAP binding site and a decrease in average mutual information at the repressor binding site. $\Delta\varepsilon_{\mathrm{pd}}$ is fixed at -5 k_BT, RNAP copy number is fixed at 1000, and repressor copy number is fixed at 10. Each data point is the mean of average mutual information across 20 synthetic datasets with the corresponding repressor binding energy. Numbered labels correspond to footprints in (B). (B) Representative information footprints where $\Delta\varepsilon_{\mathrm{rd}}$ is set to $-20\ k_BT$ and $-10\ k_BT$. (C) Increasing the copy number of the repressor leads to a decrease in average mutual information at the RNAP binding site and an increase in average mutual information at the repressor binding site. $\Delta\varepsilon_{\mathrm{pd}}$ is fixed at -5 k_BT and $\Delta\varepsilon_{\mathrm{rd}}$ is fixed at -15 k_BT. RNAP copy number is fixed at 1000. Each data point is the mean of average mutual information across 20 synthetic datasets with the corresponding repressor copy number. Numbered labels correspond to footprints in (D). (D) Representative information footprints where repressor copy numbers are set to 1 and 500.

Here, the Boltzmann weight of the repressor has been increased a hundred fold compared to Eq 27. Due to the strong binding of the repressor, mutations at the RNAP binding site do not change expression on measurable levels and therefore the signal is low at the RNAP binding site.

In particular, we see in Fig 7(A) that when the repressor energy is increased beyond $-11\ k_BT$, the average mutual information at the RNAP binding site saturates and the average mutual information at the repressor binding site remains close to 0. To explain this effect, we again take a look at the ratio between the Boltzmann weights of the repressor and RNAP, the expression for which is stated in Eq 12. Here, we fix the copy number of the repressors and RNAP, the wild-type binding energy of RNAP, the number of mutations,

and the effect of mutations. Therefore,

$$\kappa = \frac{R}{P} \cdot e^{-\beta(\Delta\varepsilon_{\mathrm{rd}} - \Delta\varepsilon_{\mathrm{pd}} + m_R \Delta\Delta\varepsilon_{\mathrm{rd}} - m_P \Delta\Delta\varepsilon_{\mathrm{pd}})} \tag{37}$$

$$= \frac{10}{1000} \cdot e^{-\beta\Delta\varepsilon_{\mathrm{rd}} - 5 - 2\times 2.24 + 2\times 0.36} \tag{38}$$

$$= 1.5 \times 10^{-6} \cdot e^{-\beta\Delta\varepsilon_{\mathrm{rd}}}. \tag{39}$$

We assume that κ needs to be at least 0.1 for there to be an observable signal at the repressor binding site. Solving for $\varepsilon_{\mathrm{rd}}$ using the above equation, we have that $\Delta\varepsilon_{\mathrm{rd}} \approx -11\ k_B T$. This matches with our observation that the signal stabilizes when $\Delta\varepsilon_{\mathrm{rd}} > -11\ k_B T$. Taken together, these result invite us to rethink our interpretation of MPRA data, as the lack of signal may not necessarily indicate the absence of binding site, but it may also be a result of weak binding or low transcription factor copy number.

2.2 Changing the regulatory logic of the promoter

In the previous section, we examined the changes in information footprints when we tune the copy number of the repressors under the simple repression regulatory architecture. The effect of transcription factor copy numbers on the information footprints is more complex when a promoter is regulated by multiple transcription factors. In particular, the changes in the information footprints depend on the regulatory logic of the promoter. To see this, we consider a promoter that is regulated by two repressors. For a double-repression promoter, there are many possible regulatory logics; two of the most common ones are an AND logic gate and an OR logic gate. As shown in Fig 8(A), if the two repressors operate under AND logic, both repressors are required to be bound for repression to occur. This may happen if each of the two repressors bind weakly at their respective binding sites but bind cooperatively with each other. On the other hand, if the two repressors operate under OR logic, then only one of the repressors is needed for repression.

We generate synthetic datasets for an AND-logic and an OR-logic double-repression promoter that are regulated by repressors R_1 and R_2. As shown in Fig 8(B) and 8(C), under AND logic, there is no signal at either of the repressor binding sites when R_1 is set to 0. This matches our expectation because AND logic dictates that when one of the two repressors is absent, the second repressor is not able to reduce the level of transcription by itself. We recover the signal at both of the repressor binding sites even when there is only one copy of R_1. Interestingly, when $R_1 > R_2$, the signal at the R_1 binding site is lower than the signal at the R_2 binding site. This may be because the higher copy number of R_2 compensates for the effects of mutations and therefore expression levels are affected to a greater extent by mutations at the R_1 binding site than mutations at the R_2 binding site. In comparison, since the two repressors act independently when they are under OR logic, the signal at the R_2 binding site is preserved even when $R_1 = 0$. Moreover, the state where R_1 represses transcription competes with the state where R_2 represses transcription. As a result, when R_1 is increased, the signal at the R_1 binding site is increased whereas the signal at the R_2 binding site decreases.

As illustrated in S8 Appendix and S9 Appendix, the procedure described above can be extended to other regulatory logic such as the XOR gate and other promoter architectures such as a promoter that is regulated by two activators. These results are informative in the context of transcription factor deletion, which is a key approach for identifying and verifying which transcription factor binds to the putative binding sites discovered in MPRA pipelines [16]. The final copy number of the transcription factor depends on which experimental method is chosen to perform the deletion. If the gene coding for the transcription factor is knocked out, no transcription factor will be expressed and the transcription factor copy number will be 0. Therefore, by comparing the footprints from the wild-type strain and the transcription factor deletion strain, we can locate the site where the signal disappears and deduce which transcription factor is bound at that site. On the other hand, if knock-down methods such as RNA interference are used, some leaky expression may take place and the transcription factor copy number may be low but non-zero. In this case, there may not be appreciable differences in the footprints from the wild-type strain and the deletion strain. This would be an important point of consideration in MPRAs where knock-down methods are used to match transcription factors to binding sites.

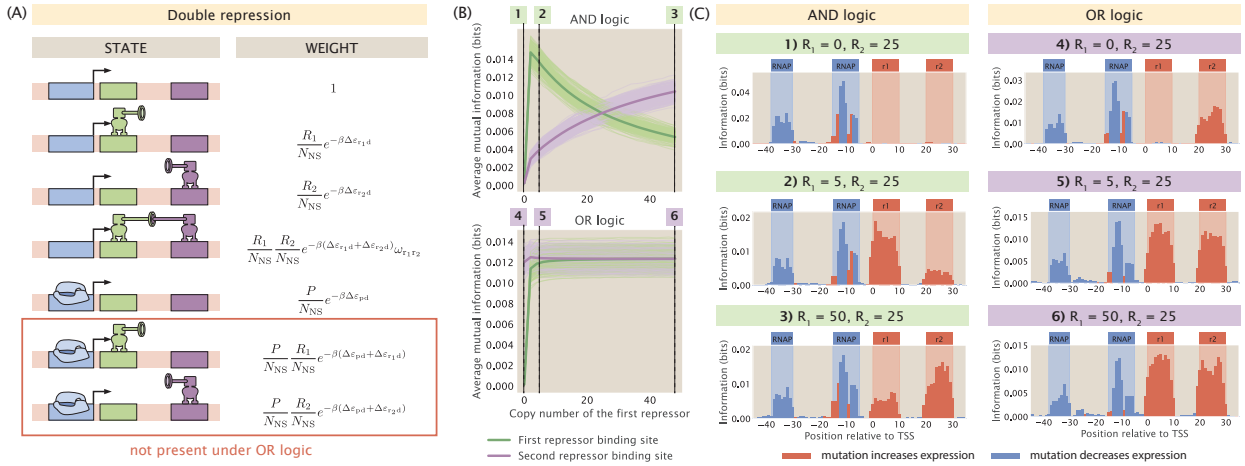

Fig 8. Changing repressor copy number for a double-repression promoter. (A) States-and-weights diagram of a promoter with the double repression regulatory architecture. The bottom two states are only present under AND logic and not present under OR logic. The states-and-weights diagram of a double repression promoter with OR logic is also shown in Fig S1(E). (B) Changing the copy number of the first repressor under AND logic and OR logic affects the signal at both repressor binding sites. For the energy matrices of the repressors, the interaction energy between the repressor and a site is set to 0 k_BT if the site has the wild-type base identity and set to 1 k_BT if the site has the mutant base identity. The interaction energy between the repressors is set to -5 k_BT. 200 synthetic datasets are simulated for each copy number. We observe that the average mutual information at binding sites has high variability across synthetic datasets, especially under OR logic. To show variability, the trajectory for each of the synthetic dataset is shown as an individual light green or light purple curve. The average trajectories across all 200 synthetic datasets are shown as the bolded green curves and the bolded purple curves. The numbered labels correspond to footprints in (C).(C) Representative information footprints of a double repression promoter under AND and OR logic.

2.3 Competition between transcription factor binding sites

Thus far, we have assumed that each transcription factor only has one specific binding site in the genome. However, many transcription factors bind to multiple promoters to regulate the transcription of different genes. For example, cyclic-AMP receptor protein (CRP), one of the most important activators in *E. coli*, regulates 330 transcription units [52]. Therefore, it is important to understand how the relationship between sequence and binding energy changes when the copy number of the transcription factor binding site is changed.

Binding site copy number is also highly relevant in the context of the experimental MPRA pipeline. When *E. coli* is the target organism, there are two main ways of delivering mutant sequences into the cell. As illustrated in Fig 9(A), the first is by directly replacing the wild type promoter with the mutant promoter using genome integration methods such as ORBIT [53, 54]. In this case, we preserve the original copy number of the binding sites. The second method is to transform the bacterial cells with plasmids carrying the promoter variant. If this approach is used, the number of transcription factor binding sites will increase by the copy number of the plasmids. We would like to understand precisely how the signal in the resulting information footprint differs between a genome integrated system and a plasmid system.

To build a synthetic dataset that involves more than one transcription factor binding site, we once again begin by building a thermodynamic model to describe the different binding events. However, in the canonical thermodynamic model that we utilized earlier, introducing multiple transcription factor binding sites would lead to a combinatorial explosion in the number of possible states. To circumvent this issue, we introduce an alternative approach based on the concept of chemical potential. Here, chemical potential corresponds to the free energy required to take an RNAP or a transcription factor out of the cellular reservoir. As shown in Fig 9(B), it is convenient to use chemical potential because in contrast to an isolated system, the resulting

model no longer imposes a constraint on the exact number of RNAP or transcription factor bound to the promoter. Instead, we can tune the chemical potential such that we constrain the average number of bound RNAPs and transcription factors. This decouples the individual binding sites and allows us to write the total partition function as a product of the partition functions at each site.

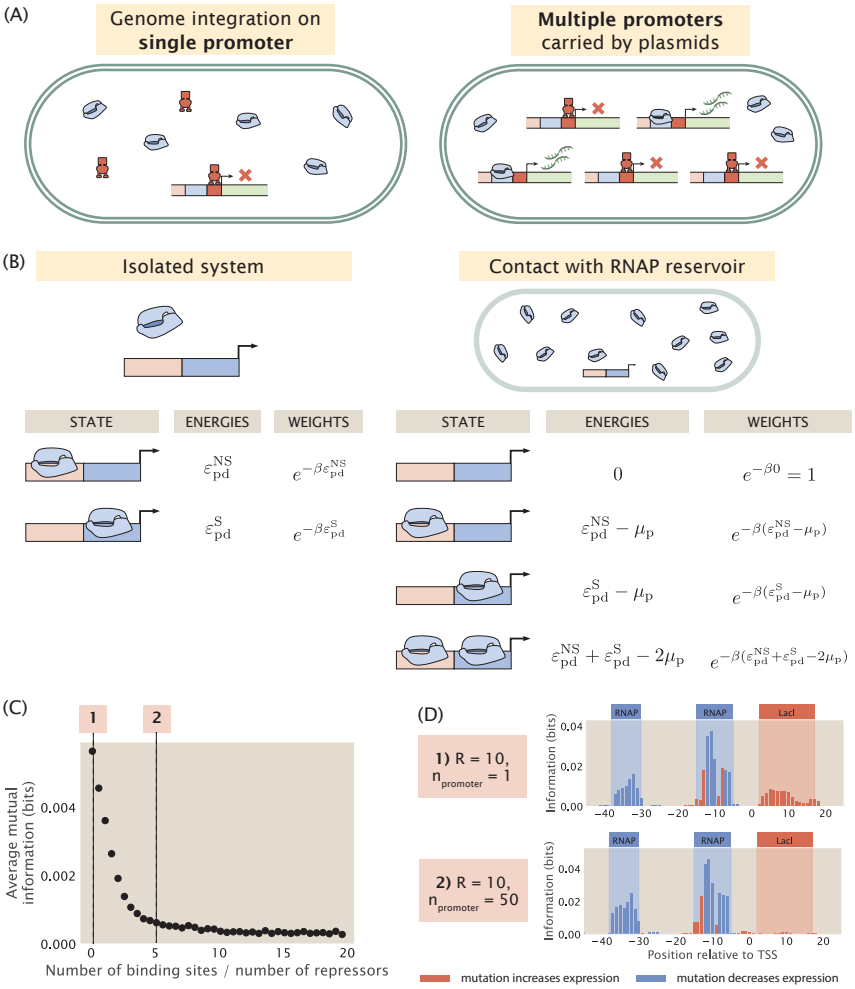

Fig 9. Changing the copy number of transcription factor binding sites. (A) There are two ways of delivering sequence variants into a cell. If the promoter variant is integrated into the genome, the original copy number of the promoter is preserved. On the other hand, if the cells are transformed with plasmids containing the promoter variant, binding site copy number will increase. When the copy number of the binding site is high, the additional binding sites titrate away the repressors and the gene will be expressed at high levels despite the presence of repressors. (B) States-and-weights models for a constitutive promoter in an isolated system and in contact with a cellular reservoir of RNAPs. To build a thermodynamic model for an isolated system, we assume that there are no free-floating RNAPs and we require that the number of bound RNAPs is equal to the total number of RNAPs in the system. On the other hand, for a system that is in contact with a reservoir, we only need to ensure that the average number of RNAPs bound matches the total number of RNAPs. (C) Average mutual information at the repressor binding site decreases when the number of the repressor binding site is increased. Repressor copy number is set to 10 for all data points. Each data point is the mean of average mutual information across 20 synthetic datasets with the corresponding number of repressor binding sites. Numbered labels correspond to footprints in (D). (D) Representative information footprints for cases where there is only 1 repressor binding site and when there are 50 repressor binding sites.

Using the method of chemical potential, we construct synthetic datasets with different repressor binding site copy numbers. As shown in Fig 9(C), as the copy number of the repressor binding site is increased, the

signal at the repressor binding site decreases rapidly and eventually stabilizes at a near-zero value. In particular, as shown in Fig 9(D), in a genome integrated system where there is only one copy of the repressor binding site, there is clear signal at the repressor binding site. On the other hand, in a plasmid system where the copy number of the binding site is greater than the copy number of the repressor, the signal for the repressor disappears. Intuitively, this is because the additional binding sites titrate away the repressors, which reduces the effective number of repressors in the system. As a result, the expression of the reporter gene no longer reflects transcriptional regulation by the repressor. This reduces the effect of mutations at the repressor binding site on expression levels, which leads to low mutual information between mutations and base identities at the repressor binding site.

In wild type *E. coli*, the median ratio of transcription factor copy number and binding site copy number is around 10 [55], and therefore the titation effects are unlikely to diminish the signals in information footprints when the sequence variant is integrated into the genome. On the other hand, if a plasmid system is used, it is beneficial to make use of a low copy number plasmid. Although we have no knowledge of which transcription factor is potentially regulating the gene of interest and therefore we do not know a priori the copy number of the transcription factor, using a low copy number plasmid has a greater chance of ensuring that the copy number of the transcription factor binding sites is no greater than the copy number of the putative transcription factor.

2.4 Changing the concentration of the inducer

So far, in the regulatory architectures that involve repressor binding, we have only considered repressors in the active state, whereas in reality the activity of the repressors can be regulated through inducer binding. Specifically, according to the Monod–Wyman–Changeux (MWC) model [**Monod1965-em**], the active and inactive states of a repressor exist in thermal equilibrium and inducer binding may shift the equilibrium in either direction. If inducer binding shifts the allosteric equilibrium of the repressor from the active state towards the inactive state, the repressor will bind more weakly to the promoter. This will increase the probability of RNAP being bound and therefore lead to higher expression. In other words, increasing inducer concentration has similar effects to knocking out the repressor from the genome. For example, when lactose is present in the absence of glucose, lactose is converted to allolactose, which acts as an inducer for the Lac repressor and leads to increased expression of genes in the LacZYA operon. Conversely, some inducer binding events may also shift the equilibrium of a repressor from the inactive state towards the active state. One example is the Trp repressor, which is activated upon tryptophan binding and represses gene expression. Here, we use the example of the lacZYA operon and demonstrate how signals in the information footprint depend on the concentration of the allolactose inducer in the system.

As shown in Fig 10(A), to include an inducible repressor in our thermodynamic model, we add an additional state in the states-and-weights diagram that accounts for binding between the inactivated repressor and the promoter. This additional state is a weak binding state where the repressor is more likely to dissociate from the binding site. In many cases, transcription factors have multiple inducer binding sites. Here, we choose a typical model where the repressor has two inducer binding sites. Based on the new states-and-weights diagram, the probability of RNAP being bound can be rewritten according to the following expression [56]

$$p_{\text{bound}} = \frac{\frac{P}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon_{\text{pd}}}}{1 + \frac{R_A}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon^A_{\text{rd}}} + \frac{R_I}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon^I_{\text{rd}}} + \frac{P}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon_{\text{pd}}}}, \tag{40}$$

where P is the number of RNAPs, R_A is the number of active repressors, R_I is the number of inactive repressors, and $\Delta\varepsilon_{\text{pd}}$, $\Delta\varepsilon^A_{\text{rd}}$, $\Delta\varepsilon^I_{\text{rd}}$ correspond to the energy differences between specific and non-specific binding of the RNAP, the active repressor, and the inactive repressor respectively.

In order to calculate the probability of RNAP being bound, we need to determine the proportion of R_A and R_I with respect to the total number of repressors. To do this, we calculate $p_{\text{active}}(c)$, the probability that the repressor exists in the active conformation as a function of the concentration of the inducer, c. To calculate $p_{\text{active}}(c)$, we model the different states of the repressor using another states-and-weights diagram,

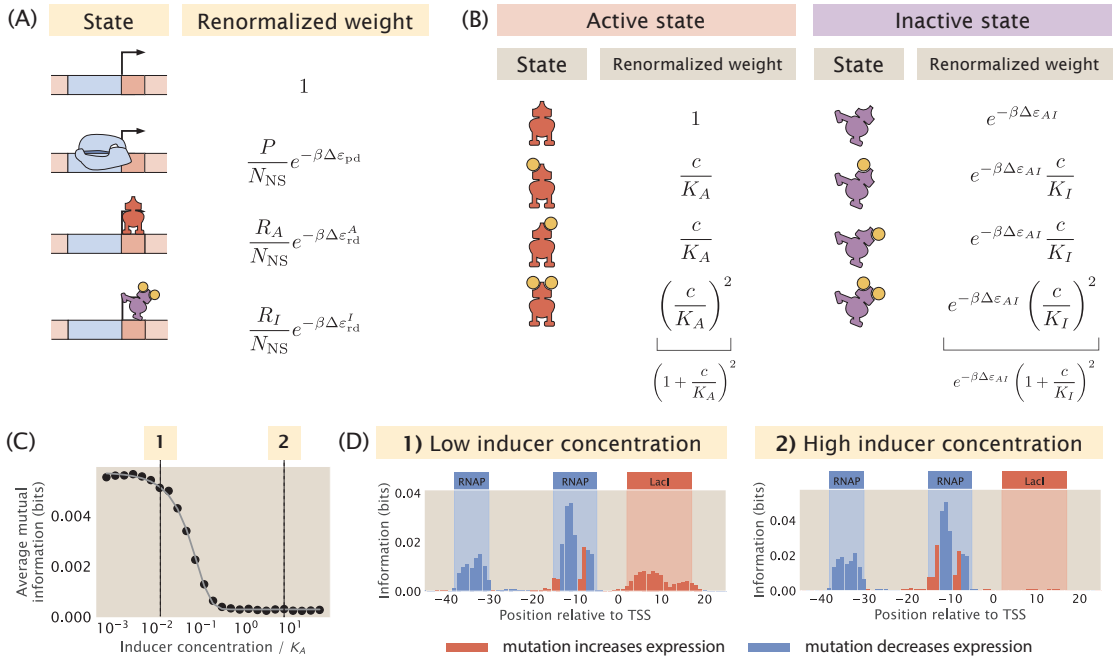

Fig 10. Changing the concentration of the inducer. (A) States-and-weights diagram for an inducible repressor. (B) States-and-weights diagram to calculate the probability that the repressor is in the active state. (C) Average mutual information at the repressor binding site decreases as the inducer concentration increases. Here, we let $K_A = 139 \times 10^{-6}$ M, $K_I = 0.53 \times 10^{-6}$ M, and $\Delta\varepsilon_{AI} = 4.5\ k_BT$. The thermodynamic parameters were inferred by Razo-Mejia et al. from predicted IPTG induction curves. [56] The inducer concentration on the x-axis is normalized with respect to the value of K_A. Each data point is the mean of average mutual information across 20 synthetic datasets with the corresponding inducer concentration. The numbered labels correspond to footprints in (D). (D) Representative information footprints with low inducer concentration (10^{-6} M) and high inducer concentration (10^{-3} M).

as illustrated in Fig 10(B). The probability that the repressor is in the active state is

$$p_{\text{active}}(c) = \frac{\left(1+\frac{c}{K_A}\right)^2}{\left(1+\frac{c}{K_A}\right)^2 + e^{-\Delta\varepsilon_{AI}/k_BT}\left(1+\frac{c}{K_I}\right)^2}, \tag{41}$$

where K_A is the dissociation constant between the inducer and the active repressor, K_I is the dissociation constant between the inducer and the inactive repressor, and $\Delta\varepsilon_{AI}$ is the structural energy difference between the active repressor and the inactive repressor. This allows us to represent the number of active and inactive repressors as $R_A = p_{\text{active}} R$ and $R_I = (1 - p_{\text{active}})R$. Therefore, our expression for p_{bound} can be modified to

$$p_{\text{bound}} = \frac{\frac{P}{N_{NS}} e^{-\beta\Delta\varepsilon_{pd}}}{1 + p_{\text{active}} \frac{R}{N_{NS}} e^{-\beta\Delta\varepsilon^A_{rd}} + (1-p_{\text{active}}) \frac{R}{N_{NS}} e^{-\beta\Delta\varepsilon^I_{rd}} + \frac{P}{N_{NS}} e^{-\beta\Delta\varepsilon_{pd}}}, \tag{42}$$

We built synthetic datasets for a promoter with the simple repression regulatory architecture with an inducible repressor. As shown in Fig 10(C) and 10(D), when the concentration of the inducer is increased from 10^{-2} K_A to 1 K_A, the average signal at the repressor binding site decreases. Interestingly, the average signal is not reduced further when the concentration is increased beyond the value of K_A. As shown in S10 Appendix, similar results are observed in the case of a simple activation promoter with an inducible activator.

These results show that the presence or absence of inducers can determine whether we will obtain a signal at the transcription factor binding site. This underlies the importance of performing experimental MPRAs under different growth conditions to ensure that we can identify binding sites that are bound by

transcription factors that are induced by specific cellular conditions. These efforts may fill in the gap in knowledge on the role of allostery in transcription, which so far has been lacking attention from studies in the field of gene regulation [57].

3 Identifying transcription factor binding sites from information footprints

3.1 Noise due to stochastic fluctuations of RNAP and transcription factor copy number

When data from MPRAs is presented in the form of information footprints, one way to annotate transcription factor binding sites is to identify regions where the signal is significantly higher than background noise. Therefore, to be able to precisely and confidently identify RNAP and transcription factor binding sites from information footprints, the footprint is required to have a sufficiently high signal-to-noise ratio. However, this may not always be the case. For example, the footprint shown in Fig 11 is obtained for the *mar* operon by Ireland et al. [16]; while the signal at the RNAP binding sites and the -20 MarR binding sites are clearly above the background noise, the signals at the Fis, MarA, and +10 MarR binding sites may easily be mistaken for noise.

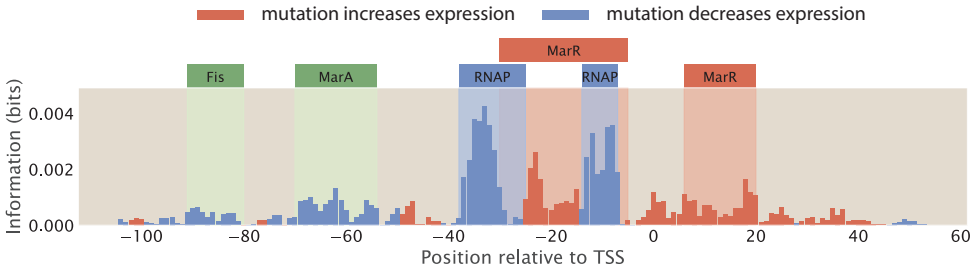

Fig 11. Annotating transcription factor binding sites by identifying sites with high signal. The footprint of the *mar* operon, produced by Ireland et al. [16]. The binding sites are annotated based on known RNAP and transcription factor binding sites; the signal at some of the binding sites, such as the Fis and MarA binding sites, are not distinguishable from background noise.

In Sec 1.3, we examined how the size of the mutant library may affect the level of noise in information footprints. Here, we continue to examine other factors that may affect signal-to-noise ratio. We first simulated possible sources of experimental noise, including PCR amplification bias and random sampling effects during RNA-Seq library preparation and RNA-Seq itself. However, as shown in S11 Appendix, these experimental processes do not lead to significant levels of noise outside of the specific binding sites. Another potential source of noise is from the biological noise that contributes to stochastic fluctuations in expression levels. These sources of biological noise can be broadly categorized into intrinsic noise and extrinsic noise [58, 59, 60]. Intrinsic noise arises from the inherent stochasticity in the process of transcription, such as changes in the rate of production or degradation of mRNA. On the other hand, extrinsic noise arises from cell-to-cell variation in the copy number of transcription machineries such as the RNAPs and transcription factors. It has been shown that extrinsic noise occurs on a longer timescale and has a greater effect on phenotypes than intrinsic noise [60]. Here, we investigate whether extrinsic noise has an effect on the information footprints.

We build synthetic datasets for a promoter with the simple repression architecture using the same procedure as before, except we no longer specify the copy number of RNAPs and repressors as a constant integer. Instead, as described in S12 Appendix, we randomly draw the copy numbers of RNAPs and repressors from a Log-Normal distribution, which is the distribution that the abundance of biomolecules often follow [61]. Brewster et al. [62] used the dilution method to measure the stochastic fluctuations in transcription factor copy numbers due to asymmetrical partitioning during cell division and they found that transcription factor copy numbers typically vary by less than the 20% of the mean copy numbers. Moreover, proteomic measurements [55, 63] suggest that the coefficient of variation for copy numbers is less than 2 even across diverse growth conditions. Here, we build Log-Normal distributions with a range of coefficients of variations, which cover both the reported levels of extrinsic noise as well as coefficients of variation as high as 100, which is physiologically extreme. As shown in Fig 12(A) and 12(B), we observe that the signal-to-noise

ratio does decrease when the extrinsic noise is higher. However, we can still distinguish between signal and noise even when we specify a large coefficient of variation. Moreover, as seen in S13 Appendix, even when the signal at the repressor binding site is low due to other factors such as low repressor binding energy, the signal from both binding sites is still distinguishable from the noise caused by copy number fluctuations. In addition, as shown in S14 Appendix, signal-to-noise ratio remains high in different regulatory architectures in the presence of extrinsic noise. These results suggest that information footprints as a read-out are robust to extrinsic noise.

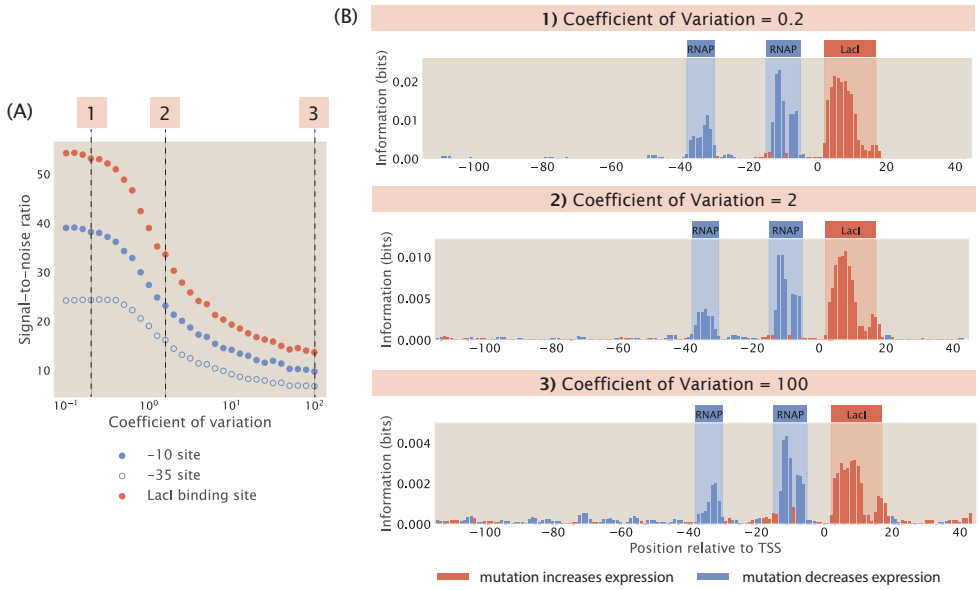

Fig 12. Adding extrinsic noise to synthetic datasets. (A) Increasing extrinsic noise lowers the signal-to-noise ratio in information footprints. For all synthetic datasets used to generate the data points, the copy numbers of RNAP and repressors are drawn using the Log-Normal distribution described in S12 Appendix. In the Log-Normal distribution, μ is set to 5000 for RNAPs and 100 for repressors. Each data point is the mean of average mutual information across 100 synthetic datasets with the corresponding coefficient of variation. The numbered labels correspond to footprints in (B). (B) Representative information footprints with three levels of extrinsic noise.

This phenomenon may be explained by the fact that the changes in binding energies due to mutations in the promoter sequence have a much greater contribution to the probability of RNAP being bound than changes in the copy number of transcription factors. Assuming that RNAP binds weakly to the promoter, the expression for p_{bound} in Eq 1 can be simplified to

$$p_{\text{bound}} = \frac{\frac{P}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon_{\text{pd}}}}{1 + \frac{R}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon_{\text{rd}}}}. \tag{43}$$

Based on the experimentally measured energy matrix for LacI [41], the average increase in $\Delta\varepsilon_{\text{rd}}$ due to one mutation is approximately 2 $k_B T$. The LacI binding site is around 20 base pairs long. Therefore, with a 10% mutation rate, there are on average 2 mutations within the LacI binding site, and the total change in $\Delta\varepsilon_{\text{rd}}$ would be approximately 4 $k_B T$. This leads to a $e^{-\beta \Delta\Delta\varepsilon_{\text{rd}}} = 0.01$ fold change in the magnitude of $\frac{R}{N_{\text{NS}}} e^{-\beta \Delta\varepsilon_{\text{rd}}}$. This means that the copy number of the repressor would have to change by a factor of 100 to overcome the effect of mutations, and this is not possible through fluctuations due to extrinsic noise. Therefore, extrinsic noise by itself will not lead to a sufficiently large change in expression levels to affect the signals in the information footprint.

3.2 Non-specific binding along the promoter

In the earlier sections of the paper, our thermodynamic models only allow RNAPs to bind at binding sites at specific positions along the promoter. However, in reality, non-specific binding events along the rest of the promoter also occur, albeit at low frequencies. To investigate the effect of non-specific binding on information footprints, we build a thermodynamic model that allow for RNAP binding at every possible position along the promoter as a unique state. The states-and-weights diagram of this expanded thermodynamic model is illustrated in Fig 13(A). The weight of each state is calculated by mapping the energy matrix to the corresponding non-specific binding site sequence at each position along the promoter. As shown in Fig 13(B) top panel, in general, non-specific binding only leads to a small amount of noise. Similar to the case of extrinsic noise, this source of noise is not at a sufficiently high level to interfere with our ability to delineate binding site positions. This implies that reducing the hitch-hiking effects described in Sec 1.3 should be the primary focus when a high signal-to-noise ratio is desired.

On the other hand, a more interesting phenomenon when RNAP is allowed to bind along the entire promoter is the presence of strong signal at non-canonical binding sites. In particular, signal may arise at these sites due to the presence of TATA-like motifs. RNAPs with σ_{70} typically bind to the TATA box, which is a motif with the consensus sequence TATAAT located at the -10 site along the promoter. However, since the sequence motif is short, it is likely that TATA-like motifs with a short mutational distance from the TATA-box sequence may exist away from the -10 site. A mutation may easily convert these motifs into a functional TATA-box, allowing RNAP to initiate transcription from a different transcription start site. In the bottom information footprint of Fig 13(B), the promoter is engineered to contain a TATA-like motif upstream of the canonical binding sites. As shown in the footprint, this leads to a strong signal at the -100 and -75 positions. This analysis unveils a feature of information footprints that may complicate interpretation of signal and noise in real-world MPRA datasets.

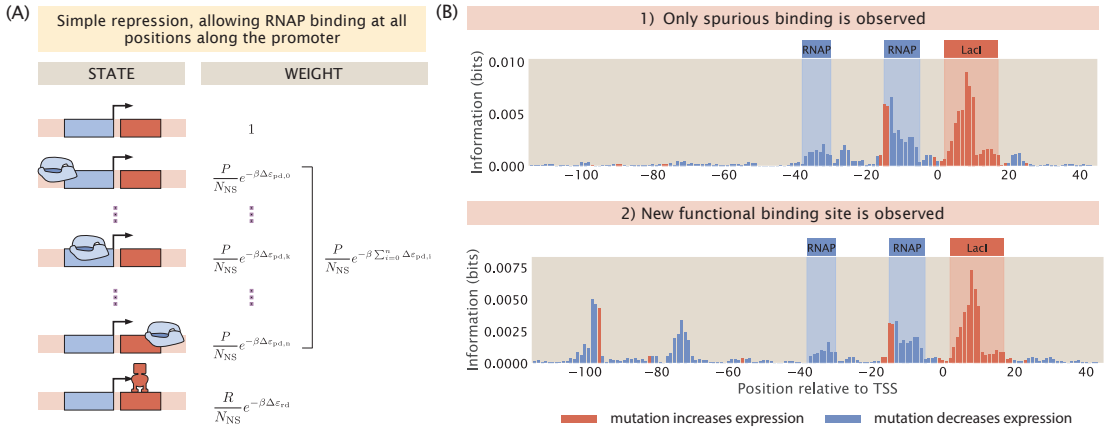

Fig 13. Non-specific RNAP binding can create low levels of noise and lead to non-canonical functional binding sites. (A) States-and-weights diagram of a simple repression promoter where spurious RNAP binding is allowed. For each of the RNAP spurious binding events, the binding energy, $\Delta\varepsilon_{\mathrm{pd,i}}$, is computed by mapping the RNAP energy matrix to the spurious binding site sequence. The index i correspond to the position of the first base pair to which RNAP binds along the promoter. 0 is at the start of the promoter sequence; k is at the canonical RNAP binding site; $n = 160 - l_p$ is index of the most downstream binding site where the promoter is assumed to be 160 base-pair long and l_p is the length of the RNAP binding site. (B) Information footprints of a promoter under the simple repression regulatory architecture with non-specific binding (top) and with a new functional binding site (bottom). The bottom plot is created by inserting the sequence "TAGAAT", which is one letter away from the TATA-box sequence, at the -80 position.

3.3 Overlapping binding sites

Other than a low signal-to-noise ratio, another factor that may contribute to the challenge of deciphering regulatory architectures is the presence of overlapping binding events. This is especially common with RNAP and repressor binding sites, since a common mechanism by which repressors act to reduce expression is by binding to the RNAP binding site and thereby sterically blocking RNAP binding. For example, in Fig 11, we can see that the MarR binding site overlaps with the RNAP binding site. Assuming perfect binding sites, a mutation in the RNAP binding site will decrease expression and a mutation in the repressor binding site will increase expression. However, if the binding sites overlap, we expect either the signals from the two binding events will cancel out or one signal will dominate the other. Here, we build synthetic datasets with different degrees of overlap between RNAP and the repressor and we examine how much overlap can be tolerated before the two binding sites are no longer distinguishable from each other.

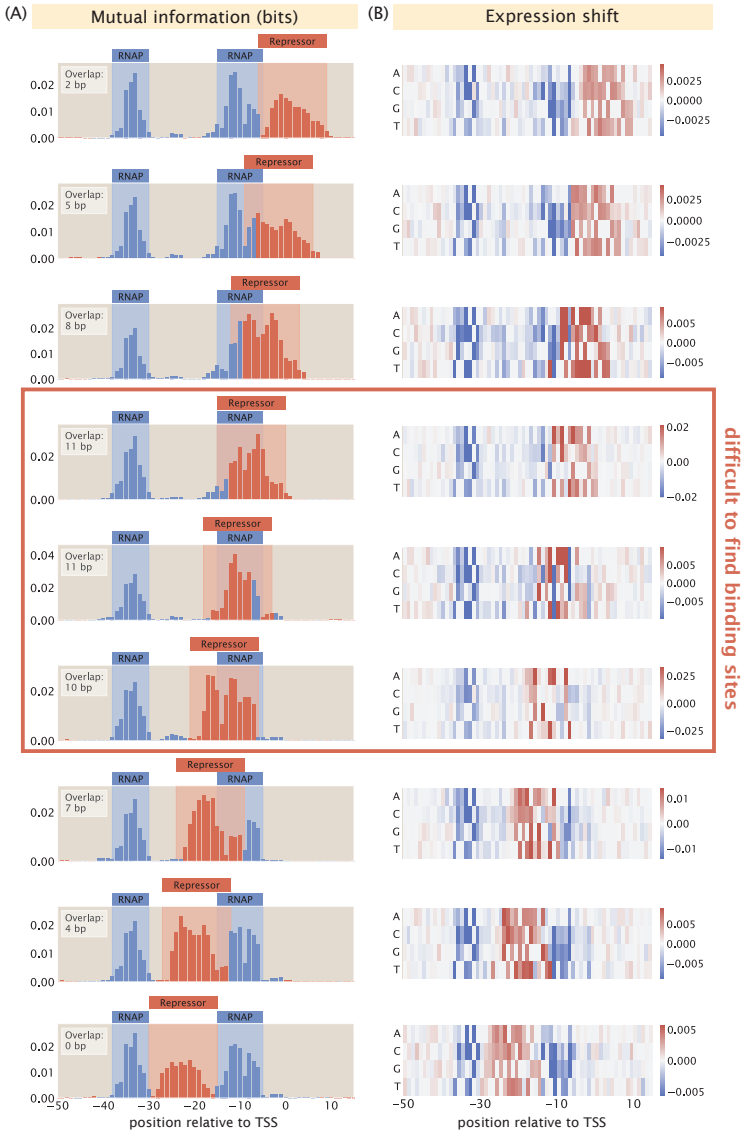

Fig 14. Changing the degree of overlap between the RNAP and repressor binding sites. (A - B) Information footprints and expression shift matrices of a simple repression promoter with overlapping binding sites. The promoters are designed to maximize binding strength given the known energy matrices of the RNAP [41] and the LacI repressor [42]. The degree of overlap in the information footprints and expression shift matrices in each row is noted at the upper left hand corner of the footprints.

In Fig 15, we show a series of information footprints and expression shift matrices where the repressor is slid along the full range of the RNAP binding site. The promoter sequences are engineered to maximize binding strengths based on the energy matrices of the RNAP and the repressor shown in Fig 2(B). At positions along the promoter that are only bound by RNAP, the base that minimizes the binding energy of the RNAP is chosen. The same applies for positions that are only bound by the repressor. On the other hand, at overlapping binding sites, base identities that minimize the total binding energy by the RNAP and the repressor are chosen. The base identities of the rest of the promoter sequence are selected at random. In the information footprints shown in Fig 14(A), the signals from the two binding events are clearly segmented and distinguishable from each other when less than 50% of the binding sites are overlapping. However, when the vast majority of the base positions are overlapping, signal from repressor binding dominates the signal from RNAP binding. In the expression shift matrices shown in Fig 14(B), when the repressor binds directly on top of the -10 RNAP binding site, some signal from RNAP binding is still preserved. However, these signals are not strong enough to highlight the presence of an RNAP binding site. Without prior knowledge that RNAP binds at this position, such a footprint could lead to the erroneous conclusion that only repressors bind to this site. These analyzes demonstrate the challenge of deciphering regulatory architectures in the presence of overlapping binding sites. This may be overcome by tuning growth conditions to reduce binding by some of the overlapping binding partners, such that we can obtain cleaner footprints with signal indicating individual binding events.

4 Building synthetic datasets under non-equilibrium conditions

So far in this paper, one important assumption underlying our thermodynamic models is that the processes involved in transcription initiation are in quasi-equilibrium. The success of thermodynamic models in predicting experimental outcomes in previous works lends credibility to the use of equilibrium models [51, 41, 62, 56, 45] in prokaryotic systems such as *E. coli*. However, in transcriptional regulation more generally, there are known to be energy consuming processes such as phosphorylation and nucleosome remodelling. Therefore, it is important to consider how our pipeline may be extended to systems where detailed balance is broken.

To account for non-equilibrium processes in the construction of synthetic MPRA datasets, we invoke the graph-theoretic approach proposed by Gunawardena [64] and used by Mahdavi, Salmon et al. [65] to calculate the probability of transcriptionally active states, which allows us to consider the full kinetic picture of trancription initiation. With the probability of transcriptionally active states calculated using this approach, we can then predict the expression levels of the promoter variants without imposing equilibrium constraints. This makes it possible for us to build a synthetic dataset that does not rely on the quasi-equilibrium assumptions, and will provide further clues about how to interpret MPRA datasets.

To see how this can be done, consider a promoter with the simple activation regulatory architecture. Recall that such a promoter can be in one of four possible states: empty (E), bound by RNAP (P), bound by the activator (A), or bound by both RNAP and the activator (AP). As shown in Fig 15(A), we can describe this architecture using a directed square graph with four vertices and eight edges. Each vertex corresponds to one of the four states; each edge describes the transition between two connected states and is associated with a rate constant. Having written down the architecture using a graph, the probability of each possible state can be derived using the Matrix Tree Theorem, which states that the probability of a given state at steady state is proportional to the sum of products of rate constants across all spanning trees that are rooted in that state. The expression for the probability of each state is given in S15 Appendix. With this, we can write down the total probability of the transcriptionally active states

$$p_{\text{active}} = p_{\text{A}} + p_{\text{AP}}, \tag{44}$$

where p_{A} is the probability that the activator is bound to the promoter and p_{AP} is the probability that both the activator and RNAP are bound to the promoter.

Importantly, to construct a synthetic MPRA dataset with broken detailed balance, the sequence-dependence of p_{active} needs to be preserved. Here, we make the simplifying assumption that all the on-rates are diffusion limited and therefore independent of promoter sequence. This means that the sequence-dependence of p_{active} comes from the mapping between the off rates and the sequences of the promoter variants. One way to create this mapping is to leverage the fact that the dissociation constant K_d

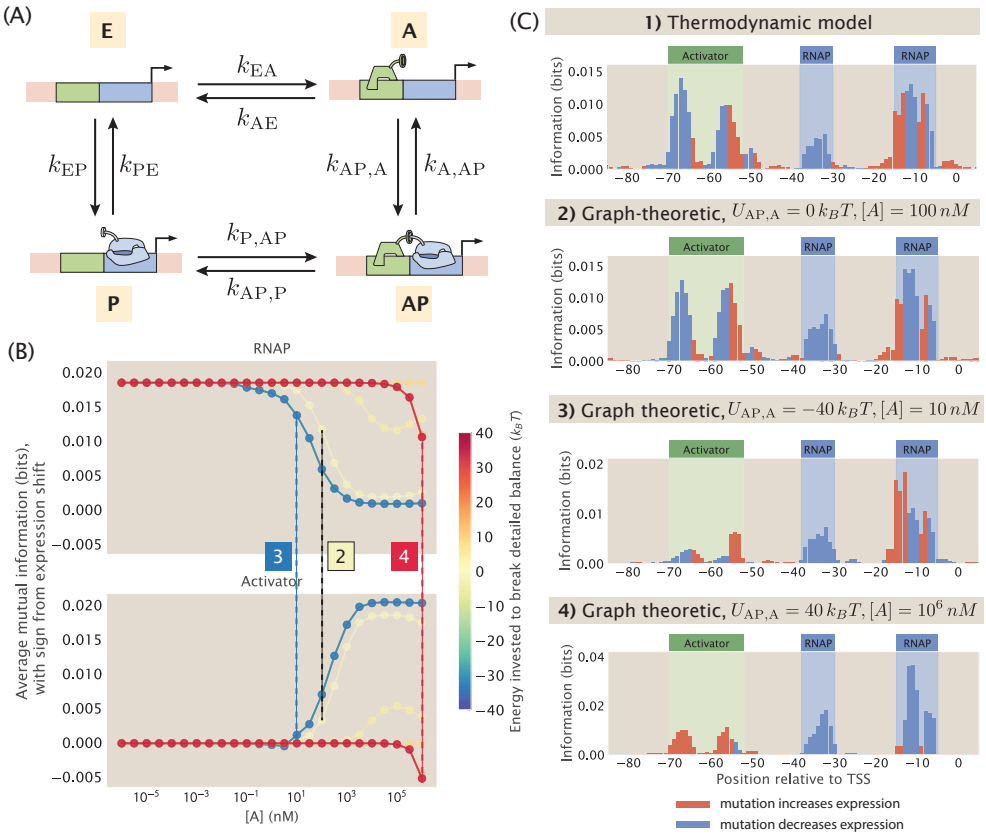

Fig 15. Building synthetic datasets with broken detailed balance. (A) Directed square graph that describes the kinetic processes of a simple activation promoter. (B) Changes in average mutual information at the RNAP and activator binding sites when the concentration of the activator ($[A]$) and the energy invested to break the detailed balance at the AP $\rightarrow$ A edge ($U_{\mathrm{AP,A}}$) are changed. The sign of the value on the y-axis is based on the expression shift values calculated using Eq 9. Each data point is the mean of average mutual information across 20 synthetic datasets with the corresponding $U_{\mathrm{AP,A}}$ and $[A]$. The numbered labels indicate datapoints for which the corresponding information footprints are shown in (C). (C) Information footprints built using the thermodynamic model and the graph-theoretic model. The second footprint is a graph-theoretic treatment of the equilibrium case.

is sequence-dependent. Specifically, we have that

$$K_d = c_0\, e^{\beta \Delta\varepsilon}, \tag{45}$$

where $c_0 = 1\ M$ is the reference concentration of the standard state. As we have demonstrated in Fig 2(A), $\Delta\varepsilon_D$ can be calculated using energy matrices in a sequence-dependent manner, which confers sequence-specificity to K_d. Since $k_{\mathrm{off}}^{\mathrm{eq}} = k_{\mathrm{on}}^{\mathrm{eq}} \times K_d$, we can also calculate k_{off} in a sequence-dependent manner. To demonstrate that the graph-theoretic approach allows us to build synthetic datasets, we first use the method outlined above to calculate k_{on} and k_{off} for each promoter variant at equilibrium and used these values to predict the expression levels of the promoter variants with the simple activation regulatory architecture. As shown in Fig 15(C), the information footprint built from the graph-theoretic synthetic dataset under equilibrium is comparable to the information footprint built from the synthetic dataset built from our thermodynamic models.

Finally, to build a synthetic dataset under non-equilibrium, we can estimate k_{off} by incorporating the energy U invested to break detailed balance, where

$$k_{\mathrm{off}} = k_{\mathrm{off}}^{\mathrm{eq}} \times e^{\beta U}. \tag{46}$$

The diffusion-limited k_{on} and the sequence-dependent k_{off} for each edge can then be used to calculate p_{active} for each promoter variant and predict the expression levels of mutant promoters under non-equilibrium conditions. For example, we can break detailed balance at the edge where RNAP unbinds from the state where both the activator and RNAP are bound. That is to say, we invest an energy $U_{\text{AP,A}}$ such that

$$k_{\text{AP,A}} = k^{\text{eq}}_{\text{AP,A}} \times e^{\beta U_{\text{AP,A}}}. \tag{47}$$

As shown in Fig 15(B), when detailed balance is broken at this edge, varying U and the concentration of the activator leads to interesting behaviour in the information footprints. Under some conditions, the footprints are similar to what is obtained under equilibrium conditions. However, as shown in panel 4 of Fig 15(C), in the cases where a positive energy is invested to break detailed balance at the AP $\rightarrow$ A edge and when activator concentration is high, the binding of activator displaces the binding of RNAP. As a result, the activator effectively behaves as a repressor. This generalized approach presents opportunities to explore MPRA datasets without being constrained by equilibrium assumptions.

Discussion

In this paper, we explore the landscape of sequence-energy-phenotype mapping by utilizing a computational pipeline that simulates MPRA pipelines. More generally, our computational pipeline makes it possible to use statistical mechanical models of gene expression to systematically explore the connection between mutations and level of gene expression. Using this pipeline, we have examined the effects of perturbing various experimental and biological parameters. These perturbations occur at multiple stages of the pipeline. Some parameters pertain to the initial library design, such as library size, mutation rate, and presence of mutational bias. Other parameters are built into the model itself, such as the copy number of the promoter and the transcription factors, which are parameters that may vary biologically or be affected by the design of experimental procedures.

We have demonstrated that our computational pipeline has high flexibility and can easily be adapted to examine the effects of other perturbations not included in this paper. Furthermore, the computational nature of the pipeline allows full parameter searches to be done precisely and efficiently. For example, it would be both time-consuming and cost-prohibitive to experimentally determine the optimal library size and mutation rate as it would involve performing a large array of experimental tests. On the other hand, using our computational pipeline, we can efficiently build a series of synthetic datasets with different mutant libraries and determine the strategy for library design that is optimal for deciphering regulatory architectures.

Apart from informing the choice of experimental parameters, the pipeline also helps to anticipate challenges involved in parsing information footprints. For example, in Sec 3.3, we predict how the signal in information footprints would be affected when there are overlapping binding sites. One potential usage of this pipeline is for building synthetic datasets that involve features that could lead to information footprints that are hard to parse. Since the synthetic datasets are built with prior knowledge of the underlying regulatory architectures, these datasets can be used to develop and improve algorithms for deciphering these architectures. This will increase confidence in the results when the same algorithms are used to analyze experimental datasets and determine the location of binding sites. Moreover, this will pave the way for automatically annotating binding sites for any given information footprint given MPRA data. To enable others who perform MPRAs in the context of transcriptional regulation to use our computational platform, we have made our code publicly available and we are developing an interactive website where the users may generate footprints given their own parameters of interest.

One limitation of our computational MPRA pipeline is that since we rely on writing down states and weights models in order to predict the probability of transcriptionally active states, the combinatorial explosion can make it challenging for us to consider promoters that are regulated by three or more transcription factors. However, data from RegulonDB would suggest that over 80% of the promoters in *E. coli* fall under the six regulatory architectures that we discussed in Fig 3 [65]. Therefore, at least in prokaryotic genomes, we are confident that our computational pipeline can be used to simulate MPRA datasets for the vast majority of promoters.

Furthermore, our current computational pipeline only considers transcription initiation factors, while other types of transcription factors, such as elongation factors and termination factors, are also important for

determining gene expression. It is challenging to include these factors into our model as it would involve additional kinetic terms that would have to be worked out, though we are extremely interested in developing these approaches as well in our future work. While this work does not directly address the challenge in understanding the role of transcription elongation factors and termination factors, we believe that by achieving a full understanding of transcription initiation factor binding through the efforts of both the computational and experimental MPRA pipelines, it will help streamline strategies needed to decipher the roles of other types of transcription factors.

In addition, our computational pipeline neglects the interaction between different genes in regulatory networks, which affects expression levels and may alter the expected signal in MPRA summary statistics such as information footprints. A future direction, therefore, involves building synthetic datasets of genetic networks. This would require an additional step where we modify the expression levels of each gene based on its dependency on other genes. This would not only improve the reliability of our prediction of expression levels, but these multi-gene synthetic datasets may also be used to test approaches for deciphering the architecture of regulatory networks.

Finally, while the vast amount of literature discussed in S1 Appendix gives us confidence on the validity of thermodynamic models, we acknowledge that there are many cases of transcriptional regulation in which detailed balance may be broken and thermodynamic models may no longer be appropriate. Our final results section (Sec 4) is a preliminary effort where we combine graph-theoretic models of transcriptional regulation and our computational pipeline to produce synthetic datasets and summary statistics without enforcing equilibrium constraints. We expect that many more interesting and informative results can come out of the angle of non-equilibrium synthetic MPRA datasets. We are excited to further pursue this direction in our future work.

In summary, we have developed a theoretical framework for a widely used category of experiments in the field of transcriptional regulation. Our simulation platform establishes a systematic way of testing how well high-throughput methods such as the MPRAs can be used to recover the ground truth of how the expression of a gene is transcriptionally regulated. This demonstrates the importance of developing theories of experiments in general, and we believe there is much untapped potential in extending similar types of theories to other areas of experimental work as well. Finally, we anticipate that this approach will also be very useful in performing systematic studies on the relation between mutations in regulatory binding sites and the corresponding level of gene expression in a way that will shed light on both physiological and evolutionary adaptation.

Data availability

All code used in this work and the presented figures are available open source at https://github.com/RPGroup-PBoC/theoretical_regseq.

Acknowledgments

We would like to thank Justin Kinney, Sara Mahdavi, and Gabriel Salmon for helpful discussions and feedback on this manuscript. This work in the Rob Phillips group is supported by the NIH Maximizing Investigators' Research Award (MIRA) 1R35 GM118043. Tom Röschinger was supported by Boehringer Ingelheim Fonds.

Author contributions

Conceptualization: Rosalind Wenshan Pan, Tom Röschinger, Rob Phillips
Formal analysis: Rosalind Wenshan Pan, Tom Röschinger
Investigation: Rosalind Wenshan Pan, Tom Röschinger, Kian Faizi
Methodology: Rosalind Wenshan Pan, Tom Röschinger, Kian Faizi
Software: Rosalind Wenshan Pan, Tom Röschinger, Kian Faizi
Visualization: Rosalind Wenshan Pan

Writing – original draft: Rosalind Wenshan Pan
Writing – review and editing: Rosalind Wenshan Pan, Tom Röschinger, Kian Faizi, Hernan Garcia, Rob Phillips
Supervision: Rob Phillips
Funding acquisition: Rob Phillips

References

1. Bartlett A, O'Malley RC, Huang SSC, Galli M, Nery JR, Gallavotti A, and Ecker JR. Mapping genome-wide transcription-factor binding sites using DAP-seq. Nat. Protoc. 2017 Aug; 12:1659–72
2. Trouillon J, Doubleday PF, and Sauer U. Genomic footprinting uncovers global transcription factor responses to amino acids in *Escherichia coli*. Cell Syst 2023 Oct; 14:860–871.e4
3. Gao Y, Lim HG, Verkler H, Szubin R, Quach D, Rodionova I, Chen K, Yurkovich JT, Cho BK, and Palsson BO. Unraveling the functions of uncharacterized transcription factors in *Escherichia coli* using ChIP-exo. Nucleic Acids Res. 2021 Sep; 49:9696–710
4. Mundade R, Ozer HG, Wei H, Prabhu L, and Lu T. Role of ChIP-seq in the discovery of transcription factor binding sites, differential gene regulation mechanism, epigenetic marks and beyond. Cell Cycle 2014; 13:2847–52
5. Bulyk ML. Computational prediction of transcription-factor binding site locations. Genome Biol. 2003 Dec; 5:201
6. Muerdter F, Boryń ŁM, and Arnold CD. STARR-seq - principles and applications. Genomics 2015 Sep; 106:145–50
7. Kinney JB and McCandlish DM. Massively Parallel Assays and Quantitative Sequence-Function Relationships. Annu. Rev. Genomics Hum. Genet. 2019 Aug; 20:99–127
8. Kinney JB, Murugan A, Callan Jr CG, and Cox EC. Using deep sequencing to characterize the biophysical mechanism of a transcriptional regulatory sequence. Proc. Natl. Acad. Sci. U. S. A. 2010 May; 107:9158–63
9. Urtecho G, Tripp AD, Insigne KD, Kim H, and Kosuri S. Systematic Dissection of Sequence Elements Controlling σ70 Promoters Using a Genomically Encoded Multiplexed Reporter Assay in *Escherichia coli*. Biochemistry 2019 Mar; 58:1539–51
10. Urtecho G, Insigne KD, Tripp AD, Brinck M, Lubock NB, Kim H, Chan T, and Kosuri S. Genome-wide Functional Characterization of *Escherichia coli* Promoters and Regulatory Elements Responsible for their Function. bioRxiv 2020 Jan :2020.01.04.894907
11. Han Y, Li W, Filko A, Li J, and Zhang F. Genome-wide promoter responses to CRISPR perturbations of regulators reveal regulatory networks in *Escherichia coli*. Nat. Commun. 2023 Sep; 14:5757
12. Vvedenskaya IO, Goldman SR, and Nickels BE. Analysis of Bacterial Transcription by "Massively Systematic Transcript End Readout," MASTER. Methods Enzymol. 2018 Oct; 612:269–302
13. Boer CG de, Vaishnav ED, Sadeh R, Abeyta EL, Friedman N, and Regev A. Deciphering eukaryotic gene-regulatory logic with 100 million random promoters. Nat. Biotechnol. 2020 Jan; 38:56–65
14. Zheng Y and VanDusen NJ. Massively Parallel Reporter Assays for High-Throughput In Vivo Analysis of Cis-Regulatory Elements. J Cardiovasc Dev Dis 2023 Mar; 10
15. Klein JC, Agarwal V, Inoue F, Keith A, Martin B, Kircher M, Ahituv N, and Shendure J. A systematic evaluation of the design and context dependencies of massively parallel reporter assays. Nat. Methods 2020 Nov; 17:1083–91
16. Ireland WT, Beeler SM, Flores-Bautista E, McCarty NS, Röschinger T, Belliveau NM, Sweredoski MJ, Moradian A, Kinney JB, and Phillips R. Deciphering the regulatory genome of *Escherichia coli*, one hundred promoters at a time. Elife 2020 Sep; 9
17. Hill A and Paganini-Hill A. The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. J. Physiol. 1910

18. Hill AV. The Combinations of Haemoglobin with Oxygen and with Carbon Monoxide. I. Biochem. J 1913 Oct; 7:471–80

19. Adair GS, Bock AV, and Field H. The hemoglobin system: VI. The oxygen dissociation curve of hemoglobin. J. Biol. Chem. 1925 Mar; 63:529–45

20. Monod J, Wyman J, and Changeux JP. On the nature of allosteric transitions: A plausible model. J. Mol. Biol. 1965 May; 12:88–118

21. Pauling L. The Oxygen Equilibrium of Hemoglobin and Its Structural Interpretation. Proc. Natl. Acad. Sci. U. S. A. 1935 Apr; 21:186–91

22. Keymer JE, Endres RG, Skoge M, Meir Y, and Wingreen NS. Chemosensing in *Escherichia coli*: two regimes of two-state receptors. Proc. Natl. Acad. Sci. U. S. A. 2006 Feb; 103:1786–91

23. Tu Y. The nonequilibrium mechanism for ultrasensitivity in a biological switch: sensing by Maxwell's demons. Proc. Natl. Acad. Sci. U. S. A. 2008 Aug; 105:11737–41

24. Mello BA and Tu Y. Quantitative modeling of sensitivity in bacterial chemotaxis: the role of coupling among different chemoreceptor species. Proc. Natl. Acad. Sci. U. S. A. 2003 Jul; 100:8223–8

25. Swem LR, Swem DL, Wingreen NS, and Bassler BL. Deducing receptor signaling parameters from in vivo analysis: LuxN/AI-1 quorum sensing in Vibrio harveyi. Cell 2008 Aug; 134:461–73

26. Buchler NE, Gerland U, and Hwa T. On schemes of combinatorial transcription logic. Proc. Natl. Acad. Sci. U. S. A. 2003 Apr; 100:5136–41

27. Kuhlman T, Zhang Z, Saier Jr MH, and Hwa T. Combinatorial transcriptional control of the lactose operon of *Escherichia coli*. Proc. Natl. Acad. Sci. U. S. A. 2007 Apr; 104:6043–8

28. Hammar P, Walldén M, Fange D, Persson F, Baltekin O, Ullman G, Leroy P, and Elf J. Direct measurement of transcription factor dissociation excludes a simple operator occupancy model for gene regulation. Nat. Genet. 2014 Apr; 46:405–8

29. Martin RG, Bartlett ES, Rosner JL, and Wall ME. Activation of the *Escherichia coli* marA/soxS/rob regulon in response to transcriptional activator concentration. J. Mol. Biol. 2008 Jul; 380:278–84

30. Ackers GK, Johnson AD, and Shea MA. Quantitative model for gene regulation by lambda phage repressor. Proc. Natl. Acad. Sci. U. S. A. 1982 Feb; 79:1129–33

31. Shea MA and Ackers GK. The OR control system of bacteriophage lambda. A physical-chemical model for gene regulation. J. Mol. Biol. 1985 Jan; 181:211–30

32. McGhee JD and Hippel PH von. Theoretical aspects of DNA-protein interactions: co-operative and non-co-operative binding of large ligands to a one-dimensional homogeneous lattice. J. Mol. Biol. 1974 Jun; 86:469–89

33. Vilar JMG and Leibler S. DNA looping and physical constraints on transcription regulation. J. Mol. Biol. 2003 Aug; 331:981–9

34. Vilar JMG and Saiz L. The unreasonable effectiveness of equilibrium gene regulation through the cell cycle. bioRxiv 2023 Apr :2023.03.31.535089

35. Kreamer NN, Phillips R, Newman DK, and Boedicker JQ. Predicting the impact of promoter variability on regulatory outputs. Sci. Rep. 2015 Dec; 5:18238

36. Eck E, Liu J, Kazemzadeh-Atoufi M, Ghoreishi S, Blythe SA, and Garcia HG. Quantitative dissection of transcription in development yields evidence for transcription-factor-driven chromatin accessibility. Elife 2020 Oct; 9

37. Meijsing SH, Pufall MA, So AY, Bates DL, Chen L, and Yamamoto KR. DNA binding site sequence directs glucocorticoid receptor structure and activity. Science 2009 Apr; 324:407–10

38. Garcia HG, Sanchez A, Boedicker JQ, Osborne M, Gelles J, Kondev J, and Phillips R. Operator sequence alters gene expression independently of transcription factor occupancy in bacteria. Cell Rep. 2012 Jul; 2:150–61

39. Bintu L, Buchler NE, Garcia HG, Gerland U, Hwa T, Kondev J, and Phillips R. Transcriptional regulation by the numbers: models. Curr. Opin. Genet. Dev. 2005 Apr; 15:116–24

40. Bintu L, Buchler NE, Garcia HG, Gerland U, Hwa T, Kondev J, Kuhlman T, and Phillips R. Transcriptional regulation by the numbers: applications. Curr. Opin. Genet. Dev. 2005 Apr; 15:125–35

41. Brewster RC, Jones DL, and Phillips R. Tuning promoter strength through RNA polymerase binding site design in *Escherichia coli*. PLoS Comput. Biol. 2012 Dec; 8:e1002811

42. Barnes SL, Belliveau NM, Ireland WT, Kinney JB, and Phillips R. Mapping DNA sequence to transcription factor binding energy in vivo. PLoS Comput. Biol. 2019 Feb; 15:e1006226

43. Bulyk ML, Johnson PLF, and Church GM. Nucleotides of transcription factor binding sites exert interdependent effects on the binding affinities of transcription factors. Nucleic Acids Res. 2002 Mar; 30:1255–61

44. Zhao Y, Ruan S, Pandey M, and Stormo GD. Improved models for transcription factor binding site identification using nonindependent interactions. Genetics 2012 Jul; 191:781–90

45. Phillips R, Belliveau NM, Chure G, Garcia HG, Razo-Mejia M, and Scholes C. Figure 1 Theory Meets Figure 2 Experiments in the Study of Gene Expression. Annu. Rev. Biophys. 2019 May; 48:121–63

46. Belliveau NM, Barnes SL, Ireland WT, Jones DL, Sweredoski MJ, Moradian A, Hess S, Kinney JB, and Phillips R. Systematic approach for dissecting the molecular mechanisms of transcriptional regulation in bacteria. Proc. Natl. Acad. Sci. U. S. A. 2018 May; 115:E4796–E4805

47. Anzalone AV, Koblan LW, and Liu DR. Genome editing with CRISPR-Cas nucleases, base editors, transposases and prime editors. Nat. Biotechnol. 2020 Jul; 38:824–44

48. Belliveau NM, Chure G, Hueschen CL, Garcia HG, Kondev J, Fisher DS, Theriot JA, and Phillips R. Fundamental limits on the rate of bacterial growth and their influence on proteomic composition. Cell Syst 2021 Sep; 12:924–944.e2

49. Bakshi S, Siryaporn A, Goulian M, and Weisshaar JC. Superresolution imaging of ribosomes and RNA polymerase in live *Escherichia coli* cells. Mol. Microbiol. 2012 Jul; 85:21–38

50. Kalisky T, Dekel E, and Alon U. Cost-benefit theory and optimal design of gene regulation functions. Phys. Biol. 2007 Nov; 4:229–45

51. Garcia HG and Phillips R. Quantitative dissection of the simple repression input-output function. Proc. Natl. Acad. Sci. U. S. A. 2011 Jul; 108:12173–8

52. Keseler IM, Gama-Castro S, Mackie A, Billington R, Bonavides-Martínez C, Caspi R, Kothari A, Krummenacker M, Midford PE, Muñiz-Rascado L, Ong WK, Paley S, Santos-Zavaleta A, Subhraveti P, Tierrafría VH, Wolfe AJ, Collado-Vides J, Paulsen IT, and Karp PD. The EcoCyc Database in 2021. Front. Microbiol. 2021 Jul; 12:711077

53. Murphy KC, Nelson SJ, Nambi S, Papavinasasundaram K, Baer CE, and Sassetti CM. ORBIT: a New Paradigm for Genetic Engineering of Mycobacterial Chromosomes. MBio 2018 Dec; 9

54. Saunders SH and Ahmed AM. ORBIT for *E. coli*: Kilobase-scale oligonucleotide recombineering at high throughput and high efficiency. bioRxiv 2023 Jun :2023.06.28.546561

55. Schmidt A, Kochanowski K, Vedelaar S, Ahrné E, Volkmer B, Callipo L, Knoops K, Bauer M, Aebersold R, and Heinemann M. The quantitative and condition-dependent *Escherichia coli* proteome. Nat. Biotechnol. 2016 Jan; 34:104–10

56. Razo-Mejia M, Barnes SL, Belliveau NM, Chure G, Einav T, Lewis M, and Phillips R. Tuning Transcriptional Regulation through Signaling: A Predictive Theory of Allosteric Induction. Cell Syst 2018 Apr; 6:456–469.e10

57. Lindsley JE and Rutter J. Whence cometh the allosterome? Proc. Natl. Acad. Sci. U. S. A. 2006 Jul; 103:10533–5

58. Elowitz MB, Levine AJ, Siggia ED, and Swain PS. Stochastic gene expression in a single cell. Science 2002 Aug; 297:1183–6

59. Fu AQ and Pachter L. Estimating intrinsic and extrinsic noise from single-cell gene expression measurements. Stat. Appl. Genet. Mol. Biol. 2016 Dec; 15:447–71

60. Raser JM and O'Shea EK. Noise in gene expression: origins, consequences, and control. Science 2005 Sep; 309:2010–3

61. Furusawa C, Suzuki T, Kashiwagi A, Yomo T, and Kaneko K. Ubiquity of log-normal distributions in intra-cellular reaction dynamics. Biophysics 2005 Apr; 1:25–31

62. Brewster RC, Weinert FM, Garcia HG, Song D, Rydenfelt M, and Phillips R. The transcription factor titration effect dictates level of gene expression. Cell 2014 Mar; 156:1312–23

63. Balakrishnan R, Mori M, Segota I, Zhang Z, Aebersold R, Ludwig C, and Hwa T. Principles of gene regulation quantitatively connect DNA to RNA and proteins in bacteria. Science 2022 Dec; 378:eabk2066

64. Gunawardena J. A linear framework for time-scale separation in nonlinear biochemical systems. PLoS One 2012 May; 7:e36321

65. Mahdavi S, Salmon GL, Daghlian P, Garcia HG, and Phillips R. Flexibility and sensitivity in gene regulation out of equilibrium. bioRxiv 2023 Apr :2023.04.11.536490