Supplemental Material

Sloppiness, robustness, and evolvability in systems biology

Bryan C. Daniels, Yan-Jiun Chen, James P. Sethna, Ryan N. Gutenkunst, and Christopher R. Myers

---

Contents

The contents of the supplemental material are organized corresponding to the order of the main text. Included in the supplemental material are derivations of mathematical results and details of the specific models mentioned in the main text.

We have also posted the data files and computer codes for the models discussed, at `http://www.lassp.cornell.edu/sethna/Sloppy`. For the KaiC, PC12, and segment polarity models, this includes:

(1) Equations in LaTeX, Python, and C
(2) SBML (system biology markup language) files
(3) Parameter ensembles
(4) Best-fit Hessian and $J^T J$, and their eigenvectors and eigenvalues

Introduction

Hessian at best fit parameters

In the introduction we mention that "the curvature of the cost surface about a best fit set of parameters is described by the Hessian H_{mn}." Examining the behavior of H_{mn} is a standard method for nonlinear least squares models when fitting data. Formally, H_{mn} is written as:

$$H_{mn} = \frac{\partial^2 C}{\partial\theta_m\theta_n} = \sum_i \frac{\partial r_i}{\partial\theta_m}\frac{\partial r_i}{\partial\theta_n} + r_i\frac{\partial^2 r_i}{\partial\theta_m\theta_n}. \tag{S1}$$

If the model fits the data well so that $r_i \approx 0$ (or perfectly, Ref. (32) in the main text) then

$$H_{mn}(\boldsymbol{\theta}^*) \approx \sum_i \frac{\partial r_i}{\partial \theta_m}\frac{\partial r_i}{\partial \theta_n} = (J^T J)_{mn}. \tag{S2}$$

If H and the cost are used (as in this review) to describe changes in model behavior from $\boldsymbol{\theta}^*$, then $\mathbf{r} \equiv 0$ at $\boldsymbol{\theta}^*$ and Equation (S2) is exact. Notice also that H reflects the sensitivity of the fit to changes in parameters; in fact, its inverse is the covariance matrix. The diagonal elements of the covariance matrix are proportional to the uncertainties in the parameters, while the off-diagonal elements are estimates of parameter uncertainty correlations.

Figure 1: Sloppiness in the mapping of chemotypes to dynatypes

Shown in Figure 1 of the main text is the mapping of chemotypes C to dynatypes D. The mapping between C and D is described with J and "J^{-1}". It is typical that $\dim(C) \ll \dim(D)$, since there are typically more data points constraining the dynatype than there are parameters defining a chemotype. Therefore, the inverse of J is not well-defined. In Figure 1, the gray ellipse in C represents the inverse image of the ϵ-ball, B_ϵ, in D under J. That is, "J^{-1}" acting on B_ϵ is the set $\{\mathbf{c} \in C \; s.t. \; J \cdot \mathbf{c} \in B_\epsilon\}$.

Note also that the stiff and sloppy eigendirections in C and their images in D can be described by the singular value decomposition of the Jacobian J. Since λ_n are eigenvalues of $J^T J$, $\sqrt{\lambda_n}$ are the singular values of J. Furthermore, writing $J = U \sum V^T$, we see that the columns of V are stiff/sloppy eigenparameters in C (shown in red in the figure), and the columns of U are images of stiff and sloppy eigenparameters (divided by λ_n) in D.

Environmental robustness and sloppiness

Figure 2: Sloppy parameter distributions: dependence on external conditions

In Figure 2 of the main text, the plane onto which the ensembles are projected is the one that aligns best with the stiffest eigenparameter of each of the four ensembles. To accomplish this, the vertical and horizontal axes in Figure 2 are, respectively, the first and second singular vectors in the singular value decomposition of the set of stiffest eigenparameters $\{\mathbf{v}_0^{25}, \mathbf{v}_0^{30}, \mathbf{v}_0^{35}, \mathbf{v}_0^{All}\}$. In a way analogous to principal components analysis, this gives the plane that passes through the origin and comes closest to passing through the heads of unit vectors pointing in the stiffest eigendirections.

Each ensemble of parameter sets shown in Figure 2 is chosen from the probability distribution corresponding to the local quadratic approximation of the cost near the best-fit parameters $\boldsymbol{\theta}^*$:

$$P(\boldsymbol{\theta}^* + \Delta\boldsymbol{\theta}) \propto \exp(-\Delta\boldsymbol{\theta} J^T J \Delta\boldsymbol{\theta}/2). \tag{S3}$$

This local approximation to the cost was used to generate the ensembles instead of the full nonlinear cost function due to difficulties in generating equilibrated ensembles: the thin curving manifolds of allowable chemotypes for sloppy models can be notoriously difficult to populate. But this is not impossible; efforts are still underway, and if equilibrated ensembles are found, they will be posted to the website mentioned above.

KaiC phosphorylation subnetwork model

In the main text, we use as an example a portion of the circadian rhythm model presented in Ref. (12) of the main text. We implement the subnetwork that van Zon et al. hypothesize must have intrinsically temperature-independent rates: that which controls the phosphorylation of KaiC alone. This subnetwork models the experimental measurements of KaiC phosphorylation in the absence of KaiA and KaiB (Ref. (11) of the main text), in which the phosphorylation of KaiC does not oscillate, but decays at a temperature-compensated rate in the range from 25 to 35° C (see circles in Figure S2).

The subnetwork involves an active and inactive state of KaiC, along with six phosphorylation sites for each state, as depicted in Figure S1. Including forward and backward "flip" rates between active and inactive states along with (de)phosphorylation rates that are each constant for the two states, there are 18 independent rates. To assess the temperature dependence, we assume that each transition rate follows an Arrhenius law, with constant energy barrier E and prefactor α: the ith rate is $\alpha_i e^{E_i/kT}$. This then gives a 36-dimensional chemotype space in which to search for solutions.

Temperature-independent solutions can be trivially found in this space if the energy barriers are chosen to be small, since this produces rates that are inherently weakly dependent on temperature. In order to avoid this trivial temperature compensation, we apply a prior that favors solutions with phosphorylation energy barriers near the expected $E_0 = 23\ kT$, similar to those found in other kinases (Ref. (46) in the main text) and appropriate for reactions that break covalent bonds. We choose this prior as a quartic in $\log E$:

$$C_{prior} = \frac{25}{2}\left[\log\left(\frac{E}{E_0}\right)\right]^4. \tag{S4}$$

The form was chosen to severely penalize barriers less than 10 kT, but to

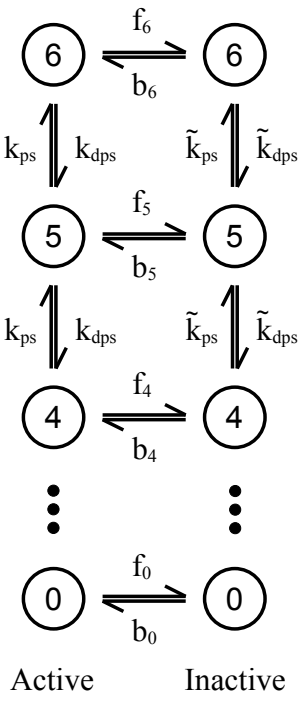

Fig. S1. **KaiC phosphorylation subnetwork.** This schematic depicts the KaiC network used as an example in the main text. It is a portion of the full circadian rhythm model presented in Ref. (12) of the main text. The numbers represent the degree of phosphorylation, and the two columns represent two different conformational states, "active" and "inactive." The labels on the arrows represent reaction rates for changing among the phosphorylation and conformation states. Each conformation state has one phosphorylation and one dephosphorylation rate, independent of the degree of phosphorylation. Each of the 14 "flip" rates between conformational states (b_i and f_i) is allowed to vary independently. This gives a total of 18 reaction rates.

be reasonably flat around E_0; other prior choices would presumably perform similarly.

Using this method, we find that it is possible to fit the experimental data even with (de)phosphorylation rates that are strongly temperature-dependent. The phosphorylation and dephosphorylation rates that provided a best fit to all temperatures simultaneously were all above 21 kT. We used Bayesian Monte-Carlo sampling of chemotype space to create an ensemble of parameter sets that each produce phosphorylation dynamics that match the experimental data at 25, 30, and 35° C. As explained above, our ensemble has not yet sampled all the space available, but we still find many such acceptable chemotypes. The minimum (de)phosphorylation rate for the ensemble was just under 10 kT, so the prior worked as designed to confine the barriers to physically reasonable values. Figure S2 shows the output of the model over this ensemble of parameter sets compared with the experimental data from Ref. (11) of the main text.

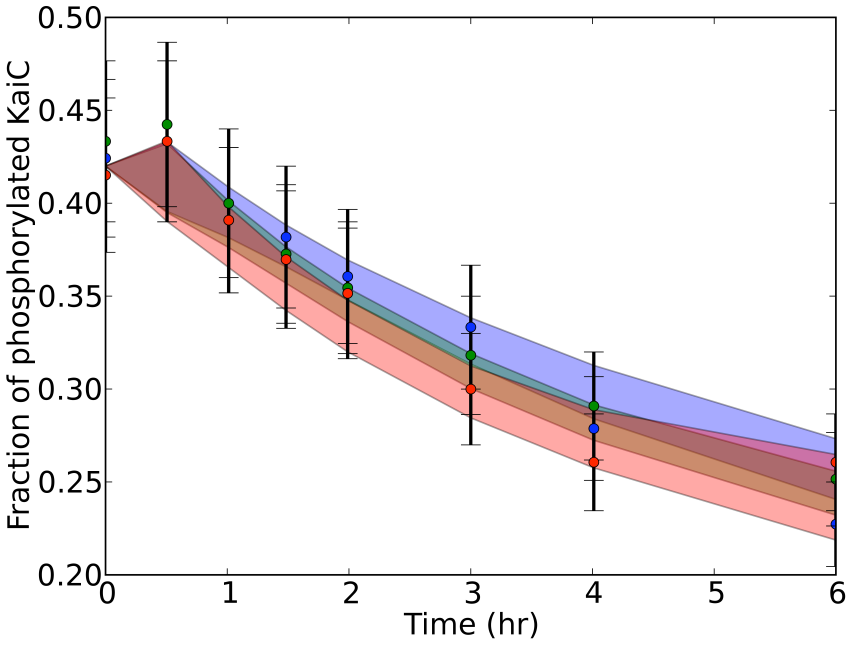

Fig. S2. **KaiC phosphorylation network: temperature-compensated output.** Shown is the net phosphorylation of KaiC over time, comparing experimental data (circles with error bars, from Ref. (11) in the main text) with output from an ensemble of chemotypes (filled colored regions, showing the mean plus or minus one standard deviation over the ensemble for the net phosphorylation at each time–point). Different colors correspond to different temperatures: blue = 25°, green = 30°, red = 35°. Note that the chemotypes describe the data well at all three temperatures, even though the rates are strongly dependent on temperature.

We mention in a footnote that, in our model, "successful chemotypes favor dephosphorylation in the active state and phosphorylation in the inactive state." This can be seen in the ratio of phosphorylation to dephosphorylation rates, shown in Figure S3, for the ensemble of successful chemotypes. Note that most members of the ensemble have an inactive state with higher phosphorylation rate than dephosphorylation, and vice versa for the active state. This matches with an intuitive temperature-compensation mechanism: with flip rates that are also temperature-dependent, higher temperatures can lead to more KaiC being in the inactive state, leading to a slower overall decay in phosphorylation that compensates for the speedup in reaction rates.

Figure 3: Sloppy model eigenvalues

The PCA shown in Figure 3 column *SP PCA* was produced after taking logarithms of the parameter values that von Dassow *et al.* used in their analysis.

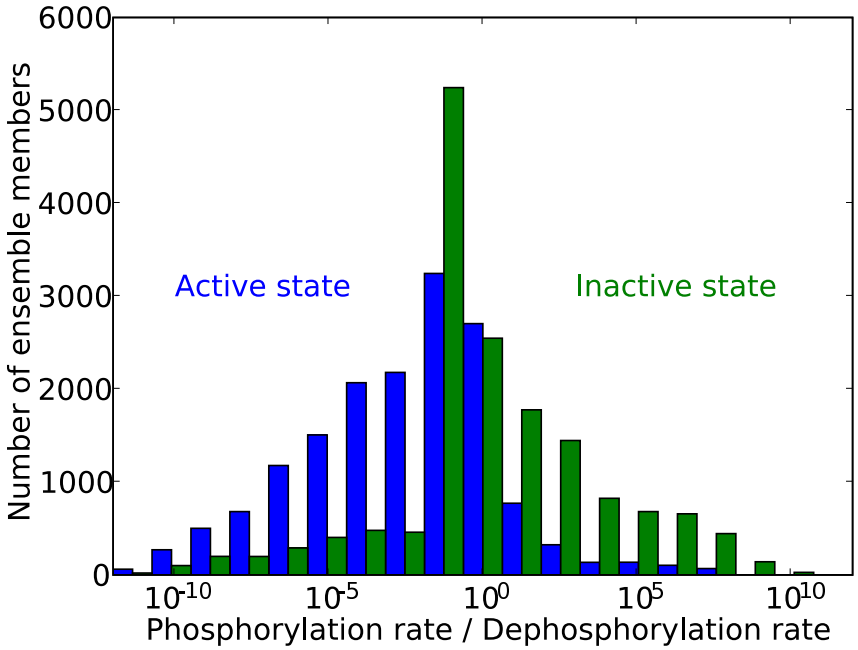

Fig. S3. **KaiC phosphorylation network: temperature-compensation mechanism.** This plot shows the ratios of phosphorylation rates to dephosphorylation rates for the active and inactive states – the distribution of k_{ps}/k_{dps} is shown in blue for the active state, and the distribution of $\tilde{k}_{ps}/\tilde{k}_{dps}$ is shown in green for the inactive state (see Figure S1 for definitions of rate constants). The distribution is over the same (non-equilibrated) ensemble as was used to generate Figure S2. Note that phosphorylation is favored in the inactive state, while dephosphorylation is favored in the active state. This suggests a temperature-compensation mechanism, as described in the text.

This measures parameter fluctuations in terms of fractional changes in parameter, rather than absolute sizes of fluctuations – allowing fluctuations in parameters with different units, for example, to be compared. The parameters used in column *SP* were chosen (logarithmic or otherwise) as defined by the original authors. Taking logarithms and/or changing units does not typically change the qualitative spectra of sloppy models, as their spectra already span so many decades.

Chemotype robustness and sloppiness

Derivation of robustness equation

In the main text (MT), the robustness is defined as

$$R_c = \prod_{\lambda_n > \lambda_{crit}} \sqrt{\frac{\lambda_{crit}}{\lambda_n}}. \tag{MT 2}$$

We now proceed to derive this result. We measure robustness as the fraction of mutations of a given size δ in C (chemotype space) that do not change the behavior beyond a given threshold (survival after a mutation), which we designate as an ϵ-ball around the optimum in D (dynatype space). Therefore we want an estimate of the fraction of the δ-ball in C that maps into the ϵ-ball in D. It is difficult to calculate this geometrically, since we would need to find the volume of an ellipsoid intersecting a sphere. Fortunately, for sloppy systems, the λ_i vary over many orders of magnitude, so we can simplify the calculation by smearing the δ-ball and ϵ-ball into Gaussians. Namely, we say a mutation $\Delta\boldsymbol{\theta}$ in C has probability $e^{-(\Delta\boldsymbol{\theta})^2/2\delta^2}/(\sqrt{2\pi}\delta)^N$, and the probability of "survival" in D is given by $e^{-r^2/2\epsilon^2}$. We then measure the robustness as the overall probability $P(\delta, \epsilon)$ of surviving after a mutation:

$$\begin{aligned} R_c &= P(\delta, \epsilon) \\ &= \left(\frac{1}{\sqrt{2\pi}\delta}\right)^N \int_C \mathrm{d}\Delta\boldsymbol{\theta}\, \exp(-(\Delta\boldsymbol{\theta})^2/2\delta^2) \exp(-(\Delta\boldsymbol{\theta})^T J^T J(\Delta\boldsymbol{\theta})/2\epsilon^2) \\ &= \prod_n \frac{1}{\sqrt{1+\lambda_n\,\delta^2/\epsilon^2}}. \end{aligned} \tag{S5}$$

For sloppy systems, λ varies over many orders of magnitude. Notice that if $\lambda_n \ll \epsilon^2/\delta^2$, its component in the product will be close to 1, and if $\lambda_n \gg \epsilon^2/\delta^2$, we can approximate the components in the product as $\sqrt{\epsilon^2/\delta^2\lambda_n}$. Therefore, using our definition $\lambda_{crit} \equiv \epsilon^2/\delta^2$ we can approximate this formula as:

$$R_c \approx \prod_{\lambda_n > \epsilon^2/\delta^2} \sqrt{\frac{\epsilon^2}{\delta^2\lambda_n}} = \prod_{\lambda_n > \lambda_{crit}} \sqrt{\frac{\lambda_{crit}}{\lambda_n}}, \tag{S6}$$

with small corrections for eigenvalues $\lambda_n \approx \epsilon^2/\delta^2$. Since this result agrees with the "slab" argument given in the main text for hard walls, we see that hard ϵ-balls and hard δ-balls will have approximately the same amount of overlap as Gaussians.

Robustness, evolvability, and sloppiness

Derivation of chemotype evolvability

In the main text, we provide a formula for the "maximum fitness change among mutations of size δ in chemotype space"

$$e_c(\mathbf{F}, \boldsymbol{\theta}) = \sqrt{\mathbf{F}^T J J^T \mathbf{F}}\delta \tag{MT 3}$$

which we derive here using a Lagrange multiplier. To derive this, we use the definition of the chemotype evolvability as the maximum response $\mathbf{r} \cdot \mathbf{F}$ in R for moves in C of size $|\Delta\boldsymbol{\theta}| = \delta$:

$$e_c(\mathbf{F}, \boldsymbol{\theta}) = \max_{|\Delta\boldsymbol{\theta}|=\delta} (\mathbf{r} \cdot \mathbf{F}). \tag{S7}$$

Next, notice that

$$\mathbf{r} \cdot \mathbf{F} = (J\Delta\boldsymbol{\theta}) \cdot \mathbf{F} = \sum_i \sum_\alpha F_i J_{i\alpha} \Delta\theta_\alpha. \tag{S8}$$

We find the optimal $\Delta\boldsymbol{\theta}$ using a Lagrange multiplier Λ. With $(\Delta\boldsymbol{\theta})^2 = \delta^2$ as our constraint, we maximize

$$F_i J_{i\alpha} \Delta\theta_\alpha + \Lambda((\Delta\boldsymbol{\theta})^2 - \delta^2) = F_i J_{i\alpha} \Delta\theta_\alpha + \Lambda(\Delta\theta_\beta \Delta\theta_\beta - \delta^2) \tag{S9}$$

where we use the Einstein summation convention (summing over repeated indices). Differentiating with respect to $\Delta\theta_\alpha$, we can find the change $\Delta\theta^{max}$ giving the maximum response:

$$\Delta\theta_\alpha^{max} = \frac{F_j J_{j\alpha}}{2\Lambda} \tag{S10}$$

and hence

$$(\Delta\boldsymbol{\theta}^{max})^2 = \frac{F_i J_{i\alpha} J_{j\alpha} F_j}{4\Lambda^2} = \delta^2, \tag{S11}$$

which implies

$$\Lambda^2 = \frac{\mathbf{F}^T J J^T \mathbf{F}}{4\delta^2}. \tag{S12}$$

Therefore, the evolvability is:

$$\begin{aligned} e_c(\mathbf{F}, \boldsymbol{\theta}) = F_i J_{i\alpha} \Delta\theta_\alpha^{max} &= \frac{F_i J_{i\alpha} F_j J_{j\alpha}}{2\Lambda} \\ &= \frac{\mathbf{F}^T J J^T \mathbf{F}}{\sqrt{\mathbf{F}^T J J^T \mathbf{F}}}\,\delta = \sqrt{\mathbf{F}^T J J^T \mathbf{F}}\,\delta. \end{aligned} \tag{S13}$$

In Equation (5), to measure overall evolvability, we defined $E_c(\boldsymbol{\theta}_\alpha)$ as a root-mean-square (RMS) average over a uniform (hyper)spherical distribution of environmental forces $\mathbf{F}$ in dynatype space. We use the RMS $\sqrt{\langle e_c(\mathbf{F},\boldsymbol{\theta}_\alpha)^2\rangle}$ rather than the average $\langle e_c(\mathbf{F},\theta_\alpha)\rangle$ because the RMS definition has an elegant result in terms of the eigenvalues λ_i of J^TJ:

$$\begin{aligned} E_c(\boldsymbol{\theta}_\alpha)^2 &= \langle e_c(\mathbf{F},\boldsymbol{\theta}_\alpha)^2\rangle_{\mathbf{F}} = \langle \mathbf{F}^T J J^T \mathbf{F}\delta^2\rangle_{\mathbf{F}} \\ &= \sum_i \frac{\int \lambda_i F_i^2 \,\mathrm{d}^N\mathbf{F}}{\int \mathrm{d}^N\mathbf{F}}\,\delta^2 \\ &= \sum_i \lambda_i \langle F_i^2\rangle\,\delta^2 = \frac{\sum_i \lambda_i \langle \mathbf{F}^2\rangle}{N}\,\delta^2 \\ &= \frac{Tr(J^TJ)\langle\mathbf{F}^2\rangle}{N}\,\delta^2 \approx \frac{Tr(H)\langle\mathbf{F}^2\rangle}{N}\delta^2. \end{aligned} \tag{S14}$$

Therefore, the overall evolvability is directly related to the trace of the Hessian:

$$E_c(\boldsymbol{\theta}_\alpha) = \sqrt{\frac{Tr(H)\langle\mathbf{F}^2\rangle}{N}}\,\delta. \tag{MT 5}$$

Our measures of robustness and evolvability depend upon our level of description, just as for Wagner's genotype and phenotype evolvabilities of RNA sequences (Ref. (8) of the main text). Our choice of an isotropic distribution of selective dynatype forces $\mathbf{F}$ is not intended as an accurate representation of actual selective forces at the phenotype level, but as an exhaustive study of all possible forces at the dynatype level of description.

Information about phenotypic selective pressures might suggest a different distribution of dynatype forces $\mathbf{F}$. Indeed, this formalism provides a mechanism for coupling maps across scales, which is an important unsolved problem. Just as the genotype-to-chemotype ($G \to C$) and chemotype-to-dynatype ($C \to D$) maps are many-to-one, so is the dynatype-to-phenotype map ($D \to P$). In the segment polarity model, for example, one might construe the phenotype as the steady-state pattern, whereas the dynatype will include information about all transient paths to that steady state. This is also closely analogous to measuring evolvability of RNA sequences by counting distinct folded structures (Ref. (8) of the main text), as many different structures may be equally nonfunctional at the higher level of biological phenotype. Ultimately, understanding the nature of the complex $D \to P$ maps will be required to estimate evolvability using more realistic distributions of selective dynatypic forces $\mathbf{F}$.

Figure 4: Evolvability and robustness in a sloppy system

When calculating the chemotype robustness R_c, we have a choice to make for the value of λ_{crit} (see Equation 2). This choice corresponds to setting the ratio of the size of acceptable changes in dynatype ϵ to the typical size of mutations δ in chemotype space: $\lambda_{crit} = \epsilon^2/\delta^2$. Equivalently, λ_{crit} sets a cutoff between stiff and sloppy eigenvalues, since we assume that, in D space, the image of the δ-ball fully overlaps with the ϵ-ball in sloppy directions (with eigenvalues below λ_{crit}), and it extends far beyond the edge of the ϵ-ball in stiff directions (with eigenvalues above λ_{crit}).

In calculating R_c for the inset of Figure 4 in the main text, we chose λ_{crit} as the fourth stiffest eigenvalue of $J^T J$ at the best fit parameters. This matches with the idea that there are only a few stiff directions that appreciably constrain parameters in chemotype space: the eigenvalues are spaced by roughly factors of three (Figure 3), meaning mutations in sloppier directions in chemotype space quickly become irrelevant in dynatype space. The choice of λ_{crit} within a reasonable range (between, say, the second stiffest and eighth stiffest eigenvalue of $J^T J$) does not qualitatively change the plot of evolvability vs. robustness.