

Emergent Predictability from Parameter Ambiguity in Biochemical Networks

Peter Eggenberger Hotz¹, Rolf Luchsinger¹, Mathias Weyland¹, Aref Enayati¹, W. David Jamieson², Oliver Castell², and Rudolf M. Fuchslin¹

¹ZHAW School of Engineering, Switzerland

²Cardiff University, United Kingdom

Corresponding author: eggg@zhaw.ch

Abstract

Complex adaptive systems often involve many interacting components and parameters, leading to strong non-identifiability: distinct parameter combinations can produce indistinguishable system behavior. Optimization methods may therefore identify parameter sets that fit observed data equally well yet differ substantially in their individual values. Despite this microscopic variability, the resulting models often retain comparable predictive power. We investigate this phenomenon in nonlinear biochemical reaction networks and show that predictive robustness arises from a low-dimensional structure in parameter space. Using Fisher Information Matrix (FIM) analysis, we identify a dominant collective parameter, Λ_1 , that acts as an effective scale controlling the system's temporal and amplitude response. Perturbations that preserve Λ_1 leave model predictions largely unchanged, whereas perturbations that alter Λ_1 substantially degrade predictive accuracy. This indicates that system behavior is governed by a low-dimensional structure embedded in a high-dimensional parameter space. These findings provide a mechanistic explanation for how robust, reproducible system-level behavior can emerge from parameter ambiguity in complex adaptive systems.

Data/Code available at: <https://doi.org/10.5281/zenodo.19485526>

Introduction

Artificial Life (ALife) seeks to understand how complex, lifelike behavior emerges from interacting components operating under local rules (Langton, 1989). A defining feature of living and lifelike systems is the coexistence of high-dimensional microscopic variability with robust, reproducible macroscopic behavior. This raises a fundamental question in Artificial Life (Bedau et al., 2000):

How can systems remain predictable when their underlying parameters are highly uncertain or even unidentifiable?

This is not merely a technical issue in model fitting. It reflects a deep biological phenomenon: cells, organisms, and

ecosystems maintain functional stability despite substantial molecular variability and environmental noise. Understanding *why* this is possible is essential both for systems biology and for the principled design of robust artificial systems.

In systems biology, this challenge appears as *parameter non-identifiability*: many different combinations of kinetic rate constants produce indistinguishable model outputs, giving rise to so-called *sloppy models* (Gutenkunst et al., 2007). While traditionally viewed as a limitation for parameter estimation, growing evidence indicates that sloppiness reflects an underlying structural property of these systems rather than a modeling artifact (Transtrum et al., 2015).

More broadly, this behavior is consistent with organizing principles observed across complex systems: macroscopic dynamics are often governed by a small number of emergent collective variables, which effectively reduce the dimensionality of the system. Examples include order parameters in statistical physics (Kadanoff, 2013), slow manifolds in dynamical systems (Transtrum et al., 2015), and coarse-grained variables in renormalization group theory (Wilson and Kogut, 1974; Kadanoff, 2013). In this view, the redundancy that renders individual parameters uncertain also enables robustness: the system can absorb microscopic variation without losing its functional identity (Kitano, 2004; Wagner, 2008).

In this work, we show that nonlinear biochemical reaction networks instantiate this principle in a concrete and quantifiable way, and we do so in a manner that is new in three respects. First, while it is known that collective fits in sloppy models yield well-constrained predictions despite large parameter uncertainty (Gutenkunst et al., 2007), and while Transtrum et al. (2010) and Transtrum et al. (2015) established the geometric structure of parameter space, the direct operational link between the dominant eigendirection and out-of-sample predictive performance has not been explicitly demonstrated. We construct a controlled perturbation experiment that shows Λ_1 is the dominant determinant of predictive accuracy within the tested perturbation range: preserving it maintains predictive accuracy across qualitatively different experimental conditions, while altering it degrades it. Second, we define $\Lambda_1 = \exp(\langle \log \theta, v_1 \rangle)$ as a concrete, computable

lumped collective variable — a weighted geometric mean of multiple rate constants — that goes beyond geometric characterization and provides a quantity that can be directly computed, monitored, and in principle targeted. Third, and most importantly for this venue, we connect this structure to foundational questions in Artificial Life. To our knowledge, the sloppy models literature, while applied across domains including physics, engineering, and power systems, has not been brought into contact with ALife questions about how coherent, robust behavior emerges from high-dimensional microscopic variability. This paper makes that connection explicit: parameter sloppiness is not merely a modeling inconvenience but a mechanism of emergent organization, directly relevant to ALife principles of robustness, evolvability, and coarse-grained description.

Our main contributions are:

- A controlled perturbation experiment demonstrating that Λ_1 is strongly associated with predictive generalization across qualitatively different experimental conditions — preserving it maintains accuracy, altering it degrades it
- The definition and operationalization of Λ_1 as a concrete, computable lumped collective variable derived from the dominant FIM eigenvector, going beyond geometric characterization to a quantity that can be directly computed and monitored
- A systematic connection of sloppy model structure to ALife principles of robustness, evolvability, and emergent organization, bridging the systems biology and ALife communities around a shared mechanism of dimensionality reduction

In short: our contribution is not the existence of sloppiness, which is known, but the identification of a concrete scalar variable Λ_1 that governs predictive equivalence across out-of-sample conditions — and the demonstration that this variable, not any individual rate constant, is a dominant determinant of predictive behavior.

Background and Related Work

Sloppy Models in Systems Biology

The concept of sloppy models was introduced by Brown and Sethna (2003) and formalized by Gutenkunst et al. (2007), who showed that the eigenvalue spectrum (i.e., the ordered set of eigenvalues) of the FIM in biological models spans many orders of magnitude. This hierarchy implies that the model is sensitive to only a few *stiff* parameter combinations while being nearly insensitive to a large number of *sloppy* directions. Subsequent work has demonstrated that sloppiness is a near-universal feature of systems biology models (Transtrum et al., 2010, 2015), and that it reflects intrinsic informational structure in the data–model relationship rather than poor model design. While collective fits in sloppy models are known to

yield well-constrained predictions (Gutenkunst et al., 2007), the present work goes further by identifying a single effective collective parameter Λ_1 , aggregating contributions from multiple rate constants, and demonstrating that preserving it maintains predictive accuracy while altering it degrades it, across distinct experimental conditions.

Emergent Low-Dimensional Structure in ALife

In Artificial Life, dimensionality reduction appears in many contexts, including self-organization, swarm behavior, neural dynamics, and evolutionary processes (e.g., Camazine et al., 2003; Couzin et al., 2002; Beer, 1995; Saxena and Cunningham, 2019; Daniels and Nemenman, 2015). Across these domains, high-dimensional microscopic interactions give rise to low-dimensional collective variables that govern macroscopic behavior. Recent work has demonstrated this principle in neural population dynamics (Saxena and Cunningham, 2019) and in the automated inference of coarse-grained dynamical descriptions from high-dimensional systems (Daniels and Nemenman, 2015).

While these studies primarily focus on state-space and behavioral dynamics, the geometry of parameter space—particularly in biochemical reaction networks—offers a complementary perspective that has received comparatively less explicit attention in ALife. The same underlying principle nevertheless applies: high-dimensional microscopic degrees of freedom collapse onto a low-dimensional effective description that governs observable behavior. What the present work adds is a concrete instantiation of this principle in parameter space rather than state space, and a quantitative test of its functional consequences for predictive robustness.

The ALife relevance is direct. A central question in ALife is how systems maintain coherent, robust behavior despite the variability and noise that are unavoidable at the microscopic level (Bedau et al., 2000; Kitano, 2004). The results presented here suggest that the geometry of parameter space may contribute to this robustness: the existence of a dominant lumped collective variable Λ_1 means the system is relatively insensitive to a vast range of microscopic configurations, as long as Λ_1 is preserved. From an evolutionary perspective, this has a further implication: systems governed by a small number of collective parameters present a vastly reduced effective search space, potentially facilitating the evolution of robust phenotypes (Wagner and Altenberg, 1996; Ciliberti et al., 2007).

The use of geometric structure in parameter space has a well-established foundation in statistical estimation theory and information geometry. The Fisher Information Matrix (FIM), originally introduced by Fisher (1925), defines a natural metric on parameter space that was later formalized geometrically by Amari and Nagaoka (2000). However, its role in explaining robustness and predictability in complex adaptive systems has not been widely explored within Artificial Life.

This paper bridges systems biology and Artificial Life by showing that parameter sloppiness corresponds to the emergence of low-dimensional collective variables that govern system behavior, framing robustness and predictability as consequences of intrinsic dimensionality reduction.

Methods

A nonlinear biochemical reaction network inspired by bioluminescent luciferase systems is used as a well-characterized testbed for complex reaction dynamics with experimentally observable outputs (Marques and Esteves da Silva, 2009), focusing on a reduced nine-state model that captures the core enzyme kinetics and luminescent emission process. Bioluminescent systems are directly relevant to synthetic and wet Artificial Life, where light-emitting gene circuits serve as canonical output signals in artificial cell research and synthetic biology. The model thus represents not merely a borrowed biological example but a system of intrinsic ALife interest.

The system is governed by mass-action kinetics:

$$\dot{E} = -k_1 EA + k_2 EAM + k_8 EP, \quad (1)$$

$$\dot{A} = -k_1 EA - k_3 AM + k_2 EAM, \quad (2)$$

$$\dot{M} = -k_3 AM + k_6 AM, \quad (3)$$

$$\dot{AM} = k_3 AM - k_6 AM - k_7 EAM, \quad (4)$$

$$E\dot{AM} = k_1 EA - k_2 EAM + k_7 EAM - k_4 EAMD, \quad (5)$$

$$\dot{D} = -k_4 EAMD + k_5 EP, \quad (6)$$

$$\dot{P} = k_4 EAMD - k_5 EP, \quad (7)$$

$$E\dot{P} = k_5 PE - k_8 EP, \quad (8)$$

$$\dot{H} = k_4 EAMD, \quad (9)$$

where products such as EA denote $E \cdot A$. The eight rate constants $k_1, \dots, k_8$ define the parameter vector θ , which is the object of all subsequent analysis.

The network follows a catalytic progression from enzyme-substrate binding through activation and light emission, with additional inhibitory branches that shape the flash profile. The resulting dynamics are strongly nonlinear and involve multiple interacting pathways, making this system a representative testbed for biochemical reaction networks relevant to synthetic and wet Artificial Life.

The parameter vector $\theta \in \mathbb{R}^8$ is estimated by minimizing a least-squares objective over time-series output data:

$$L(\theta) = \sum_i |y_i^{\text{model}}(\theta) - y_i^{\text{data}}|^2. \quad (10)$$

Synthetic training data were generated from the true parameter vector by integrating the ODE system and adding multiplicative Gaussian noise with $\sigma = 5\%$ of the instantaneous signal amplitude, producing realistic measurement variability

while preserving the qualitative shape of the luminescence flash. Fitting was performed on this noisy signal using a downsampled time grid (every 20th point of the full 6001-point series, giving 301 points per condition). The resulting SSE ($\approx 4.95 \times 10^{-18}$) is consistent with the expected noise floor for a peak signal amplitude of $\approx 1.1 \times 10^{-9}$ and 5% multiplicative noise, confirming that the optimizer is fitting the noisy data rather than the clean signal.

To explore the non-identifiable parameter space, multiple optimization runs are performed from randomized initial conditions using Differential Evolution (DE) (Storn and Price, 1997) for global exploration, followed by L-BFGS (Liu and Nocedal, 1989) for local refinement. This yields an ensemble of parameter vectors that fit the data equally well yet differ substantially in their individual values. Across eleven independent optimization runs, all converged to comparable fit quality ($\text{SSE} = 4.951 \times 10^{-18}$, $\text{CV} < 0.001\%$ across runs), while individual rate constants varied substantially across runs — k_5 alone varied 9-fold (7.6×10^3 to $6.9 \times 10^4 \text{ M}^{-1}\text{s}^{-1}$) — confirming strong non-identifiability at the level of individual parameters. The identical SSE values are not a coincidence: the sloppy model structure creates a continuous manifold of parameter combinations that all fit the data equally well, and the L-BFGS polishing step drives each run to the same residual floor set by the noise level. Different seeds therefore land at different points on this manifold, with widely varying individual parameters, but reach the same SSE.

Fitting was performed at enzyme concentration $E = 5.4 \times 10^{-6} \text{ M}$ and ATP concentration $[\text{ATP}] = 2 \times 10^{-6} \text{ M}$, the condition with the highest peak signal and therefore the greatest sensitivity to parameter variation. Fitting a single condition provides a demanding test of non-identifiability: with fewer constraints, the sloppy subspace is larger and parameter variability is maximized, making the conservation of Λ_1 more striking. Predictive generalization was assessed at three additional conditions not used during fitting: $E = 5.4 \times 10^{-6} \text{ M}$ with $[\text{ATP}] = 2 \times 10^{-7} \text{ M}$, and $E = 5.4 \times 10^{-7} \text{ M}$ with $[\text{ATP}] \in \{2 \times 10^{-6}, 2 \times 10^{-7}\} \text{ M}$, spanning qualitatively different dynamical regimes in terms of peak amplitude and timescale.

To analyze the structure of this parameter space, the Fisher Information Matrix (FIM) is computed,

$$F = J^\top J, \quad (11)$$

where $J \in \mathbb{R}^{m \times n}$ is the Jacobian of model outputs with respect to log-transformed parameters, $J_{ij} = \partial y_i / \partial \log \theta_j$. Log-transformation ensures scale invariance and reflects the multiplicative nature of biochemical rate variation.

Eigen-decomposition $F = V\Lambda V^\top$ yields principal sensitivity directions v_k ordered by decreasing eigenvalue λ_k . The leading eigenvector v_1 defines the stiffest direction in parameter space—the combination of parameters to which model output is most sensitive. The corresponding collective

parameter is defined as

$$\Lambda_1 = \exp(\langle \log \theta, v_1 \rangle), \quad (12)$$

which aggregates contributions from all rate constants according to their projection onto this dominant direction and serves as an effective collective variable capturing the system's dominant sensitivity. The exponential maps from log-parameter space back to the original scale, making Λ_1 a weighted geometric mean of rate constants — a natural measure of multiplicative variation in biochemical systems where rates span many orders of magnitude.

To test whether Λ_1 governs predictive behavior, two classes of perturbations are constructed around a reference solution θ^* . Tangential perturbations preserve Λ_1 by remaining orthogonal to v_1 , while perturbations along v_1 alter Λ_1 and leave other collective directions approximately unchanged. On-manifold perturbations were generated by randomly multiplying each rate constant (except the parameter with the largest $|v_1|$ component) by an independent factor drawn uniformly from $[1/2, 2]$ in log-space, after which that reserved parameter was adjusted analytically to restore Λ_1 to its reference value. Off-manifold perturbations were generated by scaling the same reserved parameter by fixed factors spanning $\{0.1, 0.3, 0.5, 0.8, 1.25, 2.0, 3.0, 10.0\}$, directly altering Λ_1 while leaving all other parameters unchanged. For each perturbed parameter set, the relative L_2 deviation from the reference trajectory was computed as a measure of predictive error.

Results

Table 1 shows eleven independently fitted parameter sets obtained from diverse random initial conditions (seeds spanning the range 0 to 9999999). All eleven runs converge to identical fit quality ($\text{SSE} = 4.951 \times 10^{-18}$, $\text{CV} < 0.001\%$) while individual rate constants vary substantially: k_5 varies 9-fold (7.6×10^3 to $6.9 \times 10^4 \text{ M}^{-1}\text{s}^{-1}$) and k_6 varies 2.3-fold across runs. In a generic biochemical system, a 9-fold change in k_5 — which governs the reverse reaction $EP \rightarrow P + E$ — would substantially alter equilibria, timescales, and peak amplitudes. Here it does not, because these variations remain on the sloppy subspace where the collective parameter Λ_1 is preserved — its CV across all eleven runs is just 0.004%. The three subsections below explain the geometric mechanism underlying this remarkable invariance.

A Dominant Sensitivity Direction in Parameter Space

The Fisher Information Matrix (FIM) reveals a strongly hierarchical sensitivity structure (Figure 1). The eigenvalue spectrum exhibits a single dominant eigenvalue that is separated from the remainder by many orders of magnitude: $\lambda_1/\lambda_2 \approx 1.22 \times 10^4$, spanning approximately four orders of magnitude. This indicates that model behavior is controlled primarily by one stiff direction in parameter space,




Figure 1: Eigenvalue spectrum of the Fisher Information Matrix (8 parameters, hence 8 eigenvalues). A single dominant eigenvalue is separated by multiple orders of magnitude from the remaining spectrum, indicating strong effective dimensionality reduction. The eighth (smallest) eigenvalue is numerically zero (slightly negative due to floating-point rounding) and is omitted from the log-scale plot.

while variations along all other directions have comparatively little effect. The dominance is not marginal — the second eigendirection is 1.22×10^4 times less influential than the first, making the effective dimensionality reduction unambiguous rather than approximate.

This pronounced separation implies that the high-dimensional parameter space contains a low-dimensional backbone that governs predictive behavior. In particular, the leading eigenvector defines a collective direction along which the system is highly sensitive. From an ALife perspective, this means that the effective degrees of freedom governing system behavior are far fewer than the number of microscopic parameters — a concrete instantiation of how macroscopic order emerges from microscopic complexity. The dominant eigenvector v_1 has three significant components: k_4 (+0.632), k_1 (+0.549), and k_2 (−0.547); all remaining components are below 0.001 in absolute value. Consequently, the general definition $\Lambda_1 = \exp(\langle \log \theta, v_1 \rangle)$ reduces to the approximation

$$\Lambda_1 \approx k_4^{+0.632} \cdot k_1^{+0.549} \cdot k_2^{-0.547}, \quad (13)$$

which has a direct biochemical interpretation: k_1 governs enzyme–substrate binding ($E + A \rightarrow EAM$), k_2 governs the reverse dissociation ($EAM \rightarrow E + A$), and k_4 governs product formation ($EAM \cdot D \rightarrow P$). The positive exponents on k_1 and k_4 and the negative exponent on k_2 mean that Λ_1 encodes the *net forward flux capacity* of the catalytic core: it increases when binding and product formation are fast relative to dissociation, and decreases otherwise. This geometric fact explains the pattern seen in Table 1: k_5 , which governs the reverse reaction $EP \rightarrow P + E$ and has a negligible projection onto v_1 (component < 0.001), is free to vary widely without affecting Λ_1 , while k_1 , k_2 , and k_4 , which dominate v_1 , are tightly constrained. As a check on the global stability of v_1 , the pairwise alignment angles between the leading eigenvectors computed at each of the eleven independently fitted parameter sets were all below 0.5° , confirming that the

Table 1: Eleven independently fitted parameter sets for Model 1 (random seeds 0 to 9 999 999). Bold entries mark k_5 values above 10^4 , highlighting the most variable parameter. All runs converge to identical fit quality ($\text{SSE} = 4.951 \times 10^{-18}$, $\text{CV} < 0.001\%$) while k_5 varies 9-fold (7.6×10^3 to $6.9 \times 10^4 \text{ M}^{-1}\text{s}^{-1}$). The rightmost column shows $\Lambda_1 = \exp(\langle \log \theta, v_1 \rangle)$, which remains constant to within 0.004% across all runs. Bottom row: ratio of largest to smallest value per column. Units: k_1, k_4, k_5, k_7 in $\text{M}^{-1}\text{s}^{-1}$; k_2, k_3, k_6, k_8 in s^{-1} .

Run	k_1	k_2	k_3	k_4	k_5	k_6	k_7	k_8	SSE	Λ_1
1	2.300×10^4	0.871	2.342	4821	66790	0.105	4.171×10^4	118.2	4.9508×10^{-18}	56993.0
2	2.174×10^4	0.880	2.165	5103	7624	0.240	4.285×10^4	121.4	4.9508×10^{-18}	56989.0
3	2.307×10^4	0.871	2.338	4809	61902	0.111	4.146×10^4	112.8	4.9508×10^{-18}	56991.7
4	2.205×10^4	0.877	2.348	5028	8207	0.109	4.396×10^4	75.5	4.9508×10^{-18}	56987.3
5	2.314×10^4	0.870	2.349	4789	69145	0.116	3.907×10^4	108.5	4.9508×10^{-18}	56992.2
6	2.171×10^4	0.879	2.271	5106	7603	0.230	4.965×10^4	117.2	4.9508×10^{-18}	56988.9
7	2.303×10^4	0.872	2.327	4820	64260	0.109	4.162×10^4	119.2	4.9508×10^{-18}	56993.8
8	2.300×10^4	0.870	2.278	4815	58786	0.105	4.000×10^4	99.3	4.9508×10^{-18}	56990.8
9	2.175×10^4	0.879	2.347	5099	7708	0.208	4.261×10^4	121.5	4.9508×10^{-18}	56989.3
10	2.311×10^4	0.870	2.242	4799	61796	0.114	3.851×10^4	108.0	4.9508×10^{-18}	56993.9
11	2.170×10^4	0.879	2.323	5107	7787	0.225	4.791×10^4	121.5	4.9508×10^{-18}	56989.1
Max/min	1.07	1.01	1.09	1.07	9.09	2.29	1.29	1.61	<0.001%	<0.005%

dominant sensitivity direction is stable across the biologically relevant parameter range.

Invariant Behavior Along the Low-Dimensional Manifold

The functional significance of this dominant direction is revealed by the perturbation experiments described in the Methods section. Perturbations that preserve the collective parameter Λ_1 (on-manifold) produce trajectories that closely follow the reference solution, whereas perturbations that alter Λ_1 (off-manifold) lead to large deviations (Figure 2).

Despite substantial variation in individual parameters, trajectories that preserve Λ_1 remain tightly clustered around the baseline solution. In contrast, perturbations along the dominant direction produce systematic changes in both the amplitude and temporal profile of the output. To assess the specificity of Λ_1 , note that the eigenvalue ratio $\lambda_1/\lambda_2 \approx 1.22 \times 10^4$ implies that equal-magnitude perturbations along the second eigendirection v_2 produce output deviations roughly $\sqrt{\lambda_1/\lambda_2} \approx 110$ times smaller to first order than perturbations along v_1 . The dominance of Λ_1 is therefore not a borderline result but reflects a sharp, unambiguous hierarchy in the sensitivity structure of the model. This demonstrates that predictive behavior is effectively governed by a single collective parameter — precisely the robustness mechanism ALife systems require: reliable macroscopic behavior despite microscopic variability, achieved through the geometry of parameter space rather than precise parameter tuning.

Predictive Performance Across Experimental Conditions

To assess predictive generalization, the parameter set obtained from fitting was evaluated under conditions not used




Figure 2: Model output under parameter perturbations. On-manifold perturbations (blue lines) preserve Λ_1 by varying parameters orthogonally to v_1 while analytically restoring Λ_1 to its reference value; they closely follow the reference trajectory (black). Off-manifold perturbations (red dashed) directly scale the parameter with the largest $|v_1|$ component, altering Λ_1 ; they produce large, systematic deviations. The sharp separation between on- and off-manifold trajectories demonstrates that Λ_1 is the dominant determinant of predictive accuracy within the tested perturbation range.

during estimation. Specifically, both enzyme and ATP concentrations were varied independently, spanning conditions that produce qualitatively different dynamical regimes in terms of peak amplitude and timescale.

Across all tested conditions, the fitted parameter set reproduces the observed dynamics with low error, even though the qualitative shape of the trajectories varies substantially. Prediction errors remain within the noise level of the simulated data (relative L_2 error < 5% across all tested conditions), confirming that predictive power generalizes beyond




Figure 3: Predictive performance across experimental conditions not used during fitting. Fitting was performed at $E = 5.4 \times 10^{-6} \text{ M}$, $[\text{ATP}] = 2 \times 10^{-6} \text{ M}$ (panel A). Panels B–D show predictions at three unseen conditions obtained by independently varying enzyme and ATP concentrations. Each panel shows the model prediction (red dashed) alongside clean data (green solid) and noisy observations (grey markers) with $\pm 2\sigma$ noise band. The same parameter set is used in all panels. Despite strong variations in trajectory shape and amplitude across conditions, the model accurately reproduces the observed dynamics in each case (relative L_2 error $< 5\%$), demonstrating predictive robustness beyond the fitted condition.

the fitted condition and is not limited to a single experimental configuration. Importantly, as shown in Table 1, this predictive accuracy is shared across all eleven fitted parameter sets, which show a 9-fold variation in k_5 (7.6×10^3 to $6.9 \times 10^4 \text{ M}^{-1}\text{s}^{-1}$) — a change that would substantially alter dynamics in a generic biochemical system — while Λ_1 remains constant to within 0.004%. This demonstrates that predictive equivalence is governed by Λ_1 rather than by any specific microscopic parameter configuration.

Discussion

Parameter non-identifiability has traditionally been treated as an obstacle to biological understanding: if parameters cannot be determined uniquely, what confidence can we have in model predictions? Our results reframe this question. The very structure that produces non-identifiability — the existence of a high-dimensional sloppy subspace — is also the structure that enables robustness. Systems that are non-identifiable in this structured way are *not* poorly constrained; they are constrained along a small number of meaningful directions and free along all others.

The novelty of the present contribution relative to prior work lies in making this connection operational. It was known that sloppy models yield good predictions despite parameter uncertainty (Gutenkunst et al., 2007), and that

parameter space has a geometric low-dimensional structure (Transtrum et al., 2010, 2015). What was not shown is that a single lumped collective variable Λ_1 — computable from the leading FIM eigenvector as a weighted geometric mean of multiple rate constants — is strongly associated with predictive accuracy across qualitatively different experimental regimes, within the tested perturbation range, as demonstrated by the controlled perturbation experiments reported here. Moreover, this connection has not previously been made within the Artificial Life community, where the question of how robust macroscopic behavior emerges from microscopic variability is central but has not been approached through the lens of parameter space geometry.

This reframing aligns with several currents in complexity science. The renormalization group (Wilson and Kogut, 1974; Kadanoff, 2013) formalizes how irrelevant microscopic degrees of freedom can be systematically eliminated, leaving only a small number of relevant variables that govern macroscopic behavior. Universality in statistical physics shows that macroscopic behavior can be identical across microscopically very different systems. Machta et al. (2013) have argued explicitly that sloppy models are the biological analog of renormalization group fixed points, and that parameter compression underlies emergent predictive theories. Our results are consistent with this principle: biochemical networks can be interpreted as undergoing a form of “information renormalization” in which only Λ_1 survives as a relevant variable for prediction (Machta et al., 2013; Transtrum et al., 2015). The manifold boundary approximation method (Transtrum and Qiu, 2014) offers a principled approach to extracting such reduced descriptions directly from the sloppy structure. This connection has also been explored empirically in neural population dynamics (Meshulam et al., 2019), suggesting it may be a general feature of complex biological systems.

We note that our analysis is based on local FIM geometry and therefore characterizes the sloppy structure in the vicinity of a parameter optimum. The pairwise alignment check reported in the Results section — showing alignment angles below 0.5° across all eleven independently fitted parameter sets — provides evidence that the manifold is locally flat over the biologically relevant range. Whether Λ_1 retains its predictive dominance globally across wide parameter ranges remains an open question. Geodesic methods (Transtrum et al., 2011) offer a principled approach to characterizing global manifold curvature and will be pursued in future work.

From an ALife perspective, these results offer a concrete answer to a question that has been central to the field since its inception (Langton, 1989; Bedau et al., 2000): how can complex systems maintain coherent, reproducible behavior when their microscopic components are variable, noisy, or uncertain? The geometry of parameter space itself contributes to this: a system need not be precisely tuned to behave predictably; it need only preserve its dominant collective mode Λ_1 .

Limitations and Future Work

The present study has several limitations that point toward promising directions for further investigation.

Our analysis relies on simulated data from the models themselves, which avoids experimental noise but also limits the assessment of how FIM-derived structure relates to real measurement uncertainty. Future work will apply the framework to experimental bioluminescence time series and assess whether Λ_1 identified from model analysis corresponds to identifiable experimental observables. The present demonstration is also based on fitting a single experimental condition; extending to multi-condition fitting protocols — which would further constrain the sloppy subspace and potentially reveal richer predictive structure — is a natural next step. We note that the specific parameter values reported in Table 1 depend on the particular noise realization and random seeds used; different runs will produce different individual rate constants. The key finding — that Λ_1 remains conserved to within 0.004% regardless of which parameter set is obtained — is robust across all tested seeds and noise realizations, and constitutes the central quantitative claim of the paper.

The FIM analysis is local, capturing sensitivity structure at a fixed parameter point. Since the FIM provides a local approximation, the identified direction v_1 and the associated collective parameter Λ_1 may vary across parameter space. In strongly nonlinear systems, this implies that the predictive manifold may be curved rather than flat. Geodesic methods (Transtrum et al., 2011) offer a principled approach to global manifold exploration.

Finally, the extension of this framework to other classes of ALife-relevant systems — gene regulatory networks, metabolic networks, agent-based ecological models — remains to be carried out. We expect that analogous collective parameters will emerge in these systems, but the specific geometry and its functional consequences may differ. Whether a similar principle applies to complex systems that have not been shaped by biological evolution — such as physical, social, or engineered systems — is an open question; in such systems the intrinsic coupling structure that produces sloppiness in biochemical networks may be absent or qualitatively different.

The present results also raise a longer-term question for synthetic biology and wet Artificial Life: if predictive behavior is governed by a small number of collective parameters rather than precise individual rate constants, it may be possible in principle to engineer robustness by targeting the stiff directions in parameter space rather than minimizing all parameter uncertainty. Validating this conjecture will require experimental work beyond the scope of the present study.

Conclusion

We have demonstrated that nonlinear biochemical networks can sustain accurate predictions despite severe parameter non-identifiability, and that this predictive robustness is enabled

by emergent low-dimensional structure in parameter space. A single lumped collective parameter Λ_1 , derived from the leading eigenvector of the Fisher Information Matrix and aggregating contributions from multiple rate constants, governs system behavior. Predictions are invariant under perturbations that preserve Λ_1 and sensitive to those that do not.

This finding reframes parameter sloppiness as a functional property rather than a modeling limitation. It connects the mathematical structure of biochemical kinetics to foundational ALife themes: emergent macroscopic order, developmental robustness, evolutionary evolvability, and the adequacy of coarse-grained descriptions. From an evolutionary perspective, systems governed by a small number of collective parameters present a vastly reduced search space, potentially facilitating the evolution of robust phenotypes (Wagner and Altenberg, 1996; Ciliberti et al., 2007). These results suggest that complexity need not hinder predictability—the emergence of collective variables from microscopic interaction may instead be the very mechanism by which predictability is generated (Machta et al., 2013; Transtrum et al., 2015).

Acknowledgments

This project has received funding from the European Innovation Council (EIC) under the Horizon Europe research and innovation programme, Grant Agreement No. 101194589 (Bio-HhOST).

References

- Amari, S.-i. and Nagaoka, H. (2000). *Methods of Information Geometry*. American Mathematical Society / Oxford University Press, Providence, RI. Translations of Mathematical Monographs, Vol. 191.
- Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K., and Ray, T. S. (2000). Open problems in artificial life. *Artificial Life*, 6(4):363–376.
- Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4):469–509.
- Brown, K. S. and Sethna, J. P. (2003). Statistical mechanical approaches to models with many poorly known parameters. *Physical Review E*, 68(2):021904.
- Camazine, S., Deneubourg, J.-L., Franks, N. R., Sneyd, J., Theraulaz, G., and Bonabeau, E. (2003). *Self-Organization in Biological Systems*. Princeton University Press.
- Ciliberti, S., Martin, O. C., and Wagner, A. (2007). Robustness can evolve gradually in complex regulatory gene networks. *PLoS Computational Biology*, 3(2):e15.
- Couzin, I. D., Krause, J., James, R., Ruxton, G. D., and Franks, N. R. (2002). Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218:1–11.
- Daniels, B. C. and Nemenman, I. (2015). Automated adaptive inference of phenomenological dynamical models. *Nature Communications*, 6:8133.

- Fisher, R. A. (1925). Theory of statistical estimation. *Mathematical Proceedings of the Cambridge Philosophical Society*, 22(5):700–725.
- Gutenkunst, R. N., Waterfall, J. J., Casey, F. P., Brown, K. S., Myers, C. R., and Sethna, J. P. (2007). Universally sloppy parameter sensitivities in systems biology models. *PLoS Computational Biology*, 3(10):e189.
- Kadanoff, L. P. (2013). Relating theories via renormalization. *Studies in History and Philosophy of Modern Physics*, 44:22–39.
- Kitano, H. (2004). Biological robustness. *Nature Reviews Genetics*, 5(11):826–837.
- Langton, C. G. (1989). Artificial life. In Langton, C. G., editor, *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, pages 1–47. Addison-Wesley, Redwood City, CA.
- Liu, D. C. and Nocedal, J. (1989). On the limited memory bfgs method for large scale optimization. *Mathematical Programming*, 45(1–3):503–528.
- Machta, B. B., Chachra, R., Transtrum, M. K., and Sethna, J. P. (2013). Parameter space compression underlies emergent theories and predictive models. *Science*, 342(6158):604–607.
- Marques, S. M. and Esteves da Silva, J. C. G. (2009). Firefly bioluminescence: a mechanistic approach of luciferase catalyzed reactions. *IUBMB Life*, 61(1):6–17.
- Meshulam, L., Gauthier, J. L., Brody, C. D., Tank, D. W., and Bialek, W. (2019). Coarse graining, fixed points, and scaling in a large population of neurons. *Physical Review Letters*, 123:178103.
- Saxena, S. and Cunningham, J. P. (2019). Towards the neural population doctrine. *Current Opinion in Neurobiology*, 55:103–111.
- Storn, R. and Price, K. (1997). Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4):341–359.
- Transtrum, M. K., Machta, B. B., Brown, K. S., Daniels, B. C., Myers, C. R., and Sethna, J. P. (2015). Perspective: Sloppiness and emergent theories in physics, biology, and beyond. *The Journal of Chemical Physics*, 143(1):010901.
- Transtrum, M. K., Machta, B. B., and Sethna, J. P. (2010). Why are nonlinear fits to data so challenging? *Physical Review Letters*, 104(6):060201.
- Transtrum, M. K., Machta, B. B., and Sethna, J. P. (2011). Geometry of nonlinear least squares with applications to sloppy models and optimization. *Physical Review E*, 83(3):036701.
- Transtrum, M. K. and Qiu, P. (2014). Model reduction by manifold boundaries. *Physical Review Letters*, 113(9):098701.
- Wagner, A. (2008). Robustness and evolvability: a paradox resolved. *Proceedings of the Royal Society B: Biological Sciences*, 275(1630):91–100.
- Wagner, G. P. and Altenberg, L. (1996). Perspective: Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976.
- Wilson, K. G. and Kogut, J. (1974). The renormalization group and the epsilon expansion. *Physics Reports*, 12(2):75–199.