

REVIEW

Information geometry for multiparameter models: new perspectives on the origin of simplicity

To cite this article: Katherine N Quinn *et al* 2023 *Rep. Prog. Phys.* **86** 035901

View the [article online](#) for updates and enhancements.

You may also like

- [A KIM-compliant *potfit* for fitting sloppy interatomic potentials: application to the EDIP model for silicon](#)
Mingjian Wen, Junhao Li, Peter Brommer et al.
- [Relations between entropy rate, entropy production and information geometry in linear stochastic systems](#)
Adrian-Josue Guel-Cortez and Eun-Jin Kim
- [Information geometry in vapour–liquid equilibrium](#)
Dorje C Brody and Daniel W Hook

Review

Information geometry for multiparameter models: new perspectives on the origin of simplicity

Katherine N Quinn¹ , Michael C Abbott² , Mark K Transtrum³ , Benjamin B Machta⁴ 
and James P Sethna^{5,*} 

¹ Center for the Physics of Biological Function, Princeton University, Princeton, NJ, United States of America

² Department of Physics, Yale University, New Haven, CT, United States of America

³ Department of Physics and Astronomy, Brigham Young University, Provo, UT, United States of America

⁴ Department of Physics and Systems Biology Institute, Yale University, New Haven, CT, United States of America

⁵ Department of Physics, Cornell University, Ithaca, NY, United States of America

E-mail: sethna@lassp.cornell.edu

Received 3 November 2021, revised 26 September 2022

Accepted for publication 29 November 2022

Published 28 December 2022

Corresponding editor: Dr Nandini Trivedi



Abstract

Complex models in physics, biology, economics, and engineering are often *sloppy*, meaning that the model parameters are not well determined by the model predictions for collective behavior. Many parameter combinations can vary over decades without significant changes in the predictions. This review uses information geometry to explore sloppiness and its deep relation to emergent theories. We introduce the *model manifold* of predictions, whose coordinates are the model parameters. Its *hyperribbon* structure explains why only a few parameter combinations matter for the behavior. We review recent rigorous results that connect the hierarchy of hyperribbon widths to approximation theory, and to the smoothness of model predictions under changes of the control variables. We discuss recent geodesic methods to find simpler models on nearby boundaries of the model manifold—emergent theories with fewer parameters that explain the behavior equally well. We discuss a Bayesian prior which optimizes the mutual information between model parameters and experimental data, naturally favoring points on the emergent boundary theories and thus simpler models. We introduce a ‘projected maximum likelihood’ prior that efficiently approximates this optimal prior, and contrast both to the poor behavior of the traditional Jeffreys prior. We discuss the way the renormalization group coarse-graining in statistical mechanics introduces a flow of the model manifold, and connect stiff and sloppy directions along the model manifold with relevant and irrelevant eigendirections of the renormalization group. Finally, we discuss recently developed ‘intensive’ embedding methods, allowing one to visualize the predictions of arbitrary probabilistic models as low-dimensional projections of an isometric embedding, and illustrate our method by generating the model manifold of the Ising model.

* Author to whom any correspondence should be addressed.

Keywords: information geometry, differential geometry, Bayesian, prior, hyperribbon, model reduction, emergence

(Some figures may appear in colour only in the online journal)

Contents

1. Introduction	2	a few important parameter combinations are constrained by data and are useful for making predictions, and the predictions of the model are correspondingly constrained to an effectively low-dimensional hypersurface. The fact that most complex models exhibit a hierarchy of parameter importance is known as ‘sloppiness’ [1, 2], and the hierarchical model manifold [3–6] of predictions is deemed a ‘hyperribbon’. The emergent simplicity of the model predictions, as represented by the effectively low dimensional hyperribbon, has broad implications in a variety of fields. A previous review addressed to biologists [7] considered the implications of sloppiness for evolvability and robustness of complex biological systems [8, 9], and more recent work by others has examined similar connections in neuroscience [10–12]. Writing to chemists [13], we emphasized the challenges of estimating tens to thousands of parameters from experimental data and the advantages of simple, effective theories for modeling complex systems. More broadly, there is growing need for efficient parameter reduction methods that relate detailed mechanistic models to emergent, simpler effective theories [14–20], and others have written about the connections between sloppiness, parameter identifiability, experimental design, uncertainty quantification, and other approaches to large-scale parameter estimation [21–26]. Here, we address the physics community and focus more on the geometrical study of multiparameter models and the perspective it gives on the role of parsimony in mathematical modeling.
2. The information geometry of sloppy models	3	Conventional statistical wisdom is that adding more parameters to a model is only beneficial up to a point. Eventually, additional parameters are expected to lead to worse predictions as the inferred values become sensitive to the noise rather than signal. This intuition is formalized by many results from theoretical statistics, such as the Schwartz [27] or Akaike [28] information criteria that explicitly penalize models proportionally to the number of parameters because of their capacity to over-fit. These results are formally derived in the limit of abundant, informative data, known as the asymptotic limit in theoretical statistics. But this does not match the reality of scientific inquiry in which experimental data is always limited, and in which science always contains many unobserved details. For example, systems biology effectively applies models with tens to thousands of rate and binding constants, and in machine learning, the ability of neural networks to generalize from training sets to test sets keeps growing even with millions of parameters [29]. The information geometry methods we review here provide a tangible, direct, and explicit explanation of the predictive power of models with many parameters far from asymptotic assumptions. The hierarchical nature of parameter importance in sloppy models reflects a low effective-dimensionality. This means that new parameters are typically unable to drastically improve model fits, and can therefore be
2.1. Local and global distances	5	
3. Hyperribbon bounds for nonlinear least-squares models	7	
3.1. Taylor series	7	
3.2. Chebyshev polynomials	8	
3.3. Interpolation theory	9	
4. Effective theories as manifold boundaries	10	
4.1. The manifold boundary approximation method	10	
4.2. Relation of MBAM to other reduction methods	11	
4.3. Beyond manifold boundaries	12	
5. Bayesian priors which perform well for sloppy models	13	
5.1. Maximizing mutual information	13	
5.2. Adaptive slab-and-spike priors	15	
6. Continuum limits in physics, and the renormalization group	16	
6.1. Coarse-graining and the Fisher metric	17	
6.2. Global model manifold flows	18	
7. Replicas and Minkowski embeddings for visualizing models	19	
7.1. Why Euclidean embeddings are doomed	20	
7.2. Intensive embeddings and replica theory	21	
7.3. Symmetrized Kullback–Leibler methods	21	
7.4. Intensive embedding procedure: InPCA, isKL, and MDS	22	
8. Future directions	23	
8.1. Visualization, experimental design, and testing model validity	23	
8.2. Understanding emergent low-dimensional behaviors	24	
8.3. Intensive embeddings for big data, from biology to deep neural networks	24	
8.4. Possible applications of model reduction	25	
8.5. Conclusions	26	
Data availability statement	26	
Acknowledgments	26	
References	26	

1. Introduction

Multi-parameter models in the sciences often contain many more parameters than can be accurately estimated from available data. In spite of the enormous uncertainty associated with such models, scientists regularly use them to draw physical inferences and make accurate predictions. This review summarizes recent progress toward resolving this apparent contradiction. The key observation is that in many scenarios, only

added without harming model predictions. Conversely, most parameter combinations can vary so wildly without affecting predictions, making them infeasible to measure. A reduced model description may still be preferred as being more comprehensible, reliable, or likely, but not because the fuller model will overfit.

In information geometry we view the space of all possible model predictions as forming a manifold, whose co-ordinates are the model's parameters. There is a natural notion of distance on this manifold, arising from the distinguishability of the predictions from different parameter choices. And it is an empirical fact that, for a great many models from various branches of science, this manifold has a characteristic shape we term a *hyperribbon*. This means its size in different dimensions varies widely, over many orders of magnitude, and the spectrum of these lengths is approximately evenly spaced on a log scale. The most important eigendirections need not align with our parameter directions, but their existence is the reason that simpler models can work well. Section 2 reviews these ideas, and figure 3 presents evidence that hyperribbons are ubiquitous.

Next we turn to the question of why this hyperribbon structure appears in so many systems, including those already known to admit much simpler 'spherical cow' effective models. In section 3 we show rigorously, for a widely used class of models (nonlinear least-squares models with certain smoothness approximations), that the range of possible model predictions indeed is bounded in such a way that forces our hierarchical hyperribbon structure, implying that models with tens to thousands of control knobs will nonetheless have behaviors that could be captured by a much simpler theory [30].

In many cases different simpler theories may appear in different limits. The most famous examples of this are the two reductions of 20th century physics to 19th century physics, by either turning off quantum effects, or turning off relativity. These classical theories can be viewed as living on edges of the full manifold, each with one less constant of nature, that is, one less dimension. We discuss such edges in section 4; one can generate these simpler theories for nearly any model by exploring the hierarchy of edges of the model manifold. One can remove parameter combinations by taking limits at which their specific values become unimportant, thereby obtaining a simpler model. We provide an algorithm [14] for systematically generating these reduced models using geodesics on the model manifold. For models from diverse applications ranging from power networks to systems biology [31–35], these methods have not only succeeded in generating dramatically reduced models, but the resulting models vividly illustrate key features of the complex underlying networks and match the same simplified models that experienced scientists have devised.

For Bayesian analysis of uncertainty in high-dimensional models, the volume of the manifold starts to matter. Just as high-dimensional spheres have most of their surface area near an equatorial line, so most of the volume of the model manifold is far from the edges where the simplified models dominate. The most common statistical criteria for model selection are justified in the limit of no uncertainties (excellent data). This leads to a weight (the *prior*) that is uniform over the

volume in prediction space—wildly distorting model selection in multiparameter cases by actually selecting against the simpler boundary models. In section 5 we discuss two new priors inspired by our information geometry analysis, both of which place extra weight on the manifold edges using only the anticipated uncertainty of future experimental data. On manifolds with many thin dimensions, our new priors lead to much better performance, but depend very much on the amount of data to be gathered. The first prior is obtained by maximizing mutual information [36, 37], an old idea whose consequences with finite information were long overlooked. While this is optimal, it is a discrete distribution, and not easy to solve for. The second prior (not published elsewhere) we call projected maximum likelihood, and is designed to retain many of the good properties while being much easier to sample from.

There are other principled approaches to simplifying theories, the most famous of which in physics are continuum limits and the renormalization group. These slowly discard complex microscopic behavior, to arrive at quantitatively accurate models for the long length-scale behavior—whether in condensed matter physics or quantum field theory. In section 6, we discuss the relationship of these methods to information geometry. For example, coarse-graining causes the manifold to be compressed along irrelevant dimensions, while the relevant dimensions are unchanged. This can be seen both locally, on the eigenvalues of the Fisher metric [38], and globally by studying the induced flow on the model manifold [39].

Finally, the low-dimensional emergent geometries predicted by information geometry lend themselves to direct visualization of model manifolds. For least-squares fits to data, visualization of model manifolds can be done using principal component analysis (PCA), which determines a projection into two or three dimensions based on the fraction of variation explained. Because of the hierarchical nature of the hyperribbon structure of model manifolds, PCA can effectively display the emergent low-dimensional model behavior in least-squares models. We use this method to illustrate hyperribbon bounds (figure 5) and the advantages of our new maximum likelihood and projection priors (figures 11 and 13). In section 7, we extend this visualization method using replica theory and Minkowski-like embedding spaces to address more general models that predict non-Gaussian statistical distributions. There we provide two visualizations of interest to physicists: the two-dimensional Ising model from statistical mechanics and the cosmological Λ CDM model for the fluctuations in the cosmic microwave background radiation.

2. The information geometry of sloppy models

Information geometry is where differential geometry meets statistics and information science. Models and data have geometric structures which reveal underlying relationships. For instance, the models discussed in section 4 form high-dimensional manifolds that can be embedded in a space of possible model predictions. Their connections to one another are revealed in the way their manifolds relate to each other. Furthermore, the phenomena of ill-posed and well-constrained




Figure 1. Model manifolds for a toy nonlinear model, the sum of two exponential curves with different decay rates $\theta_1, \theta_2 \geq 0$ (from equation (2)). Model manifolds are constructed as the set of all possible predictions for all possible model parameters and seen as a mapping from parameter space to prediction space. Fitting model parameters to data can be interpreted as projecting data (the red star) onto the manifold.

parameter combinations (sloppy and stiff) manifest geometrically as short and long directions along the model manifold.

The simplest scenario for a geometric interpretation are *nonlinear least-squares* models which, given input parameters, produce a prediction or output as a high-dimensional vector. Often these involve time: The model is some nonlinear function $y_\theta(t)$ with D model parameters $\theta = \{\theta^\mu\}$ evaluated at points (t_1, t_2, \dots) . We call the set of all possible model predictions for all possible model parameters the model manifold \mathcal{Y} :

$$\mathcal{Y} = \{Y(\theta) | \theta\} \quad \text{where} \quad Y(\theta) = [y_\theta(t_1), y_\theta(t_2), \dots, y_\theta(t_M)]. \quad (1)$$

The intrinsic dimension of the manifold is (at most) the number of input parameters, D . We describe it as embedded in a space whose dimension, M , is determined by the number of predictions made. If the model predicts one number $y_\theta(t_i)$ at each time t_i , then this construction embeds the manifold in Euclidean space, $\mathcal{Y} \subset \mathbb{R}^M$. (The general definition for the model manifold is given in section 2.1.)

As an illustration of this, consider a classic nonlinear model of decaying exponentials with two input parameters θ_1 and θ_2 :

$$y_\theta(t) = e^{-\theta_1 t} + e^{-\theta_2 t}, \quad \theta_\mu \geq 0. \quad (2)$$

Suppose that we evaluate this at three points in time, (t_1, t_2, t_3) . Then as we sweep through the different parameter values, we generate a curved 2D manifold embedded in a flat 3D space, as shown in figure 1. The geometry of this manifold reflects properties of the model itself: there is a symmetry to the model, since θ_1 and θ_2 can be interchanged without changing the prediction, and these points are identified in the manifold, since it folds on itself about $\theta_1 = \theta_2$. The boundaries represent important limits of the model, such as when θ_1 or θ_2 go to infinity (and thus the corresponding exponential goes to zero), or when $\theta_1 = \theta_2$. Note that boundaries are lines, 1D sub-manifolds,

corresponding to one fewer degree of freedom. These represent the simpler models obtained by taking the corresponding limit⁶.

Fitting model parameters to data has an elegant interpretation in this framework: it can be seen as projecting data onto the model manifold, as illustrated in figure 1. The red star is the measured data (x_1, x_2, \dots, x_M) , viewed as a point in the embedding space of the model manifold. To find the parameters at which the vector of predictions $[y_\theta(t_1), y_\theta(t_2), \dots, y_\theta(t_M)]$ is closest to the data point, we should minimize a sum of squares⁷,

$$\chi^2 = \sum_{i=1}^M \frac{(y_\theta(t_i) - x_i)^2}{2\sigma_i^2}. \quad (3)$$

Minimizing this measure is identical to choosing the maximum likelihood parameter, for a probabilistic model in which the observed data point x is the prediction y corrupted with Gaussian noise. This may be written $x = y + \mathcal{N}(0, \sigma^2)$, or more explicitly:

$$p(x|\theta) = \frac{1}{(2\pi\sigma^2)^{M/2}} \exp\left(-\frac{\sum_i (x_i - y_\theta(t_i))^2}{2\sigma^2}\right). \quad (4)$$

Then $\chi^2 = -\log p(x|\theta)$ is precisely the cost being minimized, (equation (3)), up to an irrelevant additive constant.

In this interpretation, the fitting of model parameters to data always involves assumptions about the noise. None of our measurements are perfectly precise, and ignoring this just means building in unstated assumptions. Instead, we interpret

⁶ Our ‘model manifold’ is not strictly a manifold. The interior of the model manifold fits the definition: every neighborhood has D co-ordinates in the usual way. But the important simpler models at the boundary introduce fold lines, corners, and other singularities that violate the formal mathematical definition of a manifold with boundaries.

⁷ Hence the name ‘nonlinear least-squares’ model.

every model as a probabilistic model for what data we will see, x . Given an outcome y as a deterministic function of the inputs, θ and t_i , we explicitly incorporate the experimental noise to generate a predicted probability distribution $p(x|\theta)$ around y . This allows us to extend what we have said about least-squares model manifolds to apply to more general models.

We now turn to how we measure distance on the model manifold. The embedding $\mathcal{Y} \subset \mathbb{R}^M$ described above is special because it is *isometric*: the Euclidean \mathbb{R}^M distances between points in \mathcal{Y} locally agree with the natural measure of statistical distance, the Fisher metric.

2.1. Local and global distances

To judge whether different parameter values are similar, or points on the model manifold are nearby, we need some notion of distance. The idea of information geometry [3–6] is that this must come from the model itself, and the noise in our measurements: if two different values predict data which is easily distinguishable, then we should not regard them as being near to each other. Conversely, if two different parameter values are close together, they must predict nearly indistinguishable data.

For nonlinear least-squares models, the distance between the predictions of two parameters θ and $\tilde{\theta}$ is given by the χ^2 distance of equation (3), $\sum_{i=1}^M (y_\theta(t_i) - y_{\tilde{\theta}}(t_i))^2 / 2\sigma_i^2$, which is basically the Euclidean distance between the two vectors y_θ and $y_{\tilde{\theta}}$. For more general models, distinguishability is measured between the probability distributions $p(x|\theta)$ predicted for a set of possible outcomes x (spin configurations for the Ising model [40], temperature maps for the cosmic background radiation [41], etc). We shall discuss three such measures, valid for general probability distributions, in section 7: the Hellinger distance [42] (equation (17)), the Bhattacharyya divergence [43] (equation (19)), and the Jeffreys [44, p 158] or symmetrized Kullback–Leibler [45, equation 2.5] divergence (equation (21)).

All of these measures are zero when $\theta = \tilde{\theta}$, are symmetric, and can be calculated for any two points θ and $\tilde{\theta}$; they are thus a global measure of distance. All of these measures reduce to the χ^2 Euclidean distance for the vector predictions of least-squares models. The latter two measures do not obey the triangle inequality, but each corresponds to distances in a Minkowski-like space where some directions are ‘time-like’ (section 7). These three, together with a wide family of other divergences [46, 47], share the same local measure, that is, the distance over an infinitesimal change $\tilde{\theta} = \theta + d\theta$:

$$ds^2 = \sum_{\mu, \nu=1}^D g_{\mu\nu} d\theta^\mu d\theta^\nu, \quad g_{\mu\nu}(\theta) = \sum_x \frac{\partial \log p(x|\theta)}{\partial \theta^\mu} \frac{\partial \log p(x|\theta)}{\partial \theta^\nu} p(x|\theta). \quad (5)$$

This is the Fisher metric, or $g_{\mu\nu}$ for $\mu, \nu = 1, 2, \dots, D$ the *Fisher information matrix* (FIM) [48]. It may be thought of as measuring distance in parameter space in units of standard




Figure 2. The prediction/behavior space of complex systems is primarily controlled by a few parameter combinations, even when the model has numerous parameters. The hierarchy of parameter importance directly translates to a geometric hierarchy in the local properties (the eigenvalues of the metric roughly form a geometric series) and global features (the manifold widths have a similar hierarchical decay). Reproduced with permission from [52]. © Copyright 2017 Annual Reviews.

deviations, using the width of the distribution $p(x|\theta)$ on the space of possible data. (Indeed, the famous Cramér–Rao bound [49–51] uses the Fisher information metric to bound the ability to distinguish the predictions between two nearby models—the ability to know that the parameters $\tilde{\theta}$ are incorrect given data generated by parameters θ .) We shall study these more general model manifolds for probabilistic models by generating *isometric embeddings*. By preserving one of the distance measures $d^2(\theta, \tilde{\theta})$ between every pair of models θ and $\tilde{\theta}$, we generate a geometrical representation of the model manifold that both preserves the FIM and (in the cases studied so far) exhibits the same sloppy hyperribbon structure as we have found in least-squares models.

While individual components of the FIM tell us how sensitive the model is to particular parameters, we are more interested in its sensitivity to parameter combinations. In the extreme case, many different parameters may control the same mechanism, making the model sensitive to changes in each of them in isolation. But if they all have equivalent effects, then the model is effectively one-dimensional, and it will be possible to make correlated changes which cancel, and do not affect the results. In this case, we would expect the FIM to have one large eigenvalue, and many small or zero ones. The one important eigenvector we will call the *stiff* direction in parameter space, and the remaining unimportant ones *sloppy*. The stiff direction is both important for predictions, and tightly constrained by observed data; the sloppy directions are neither. Figure 2 draws this in the context of an embedding $\mathcal{Y} \subset \mathbb{R}^3$.

In more realistic models with many parameters, the eigenvalues of the FIM often have a characteristic spectrum: they are spread over many orders of magnitude, and are roughly evenly spaced on a log scale. Figure 3 shows some examples of this, from explicit computation of the FIM for models from a wide variety of sciences. It is a striking empirical fact that so many of them display this structure. These eigenvalues tell us about the local geometry of the model manifold, that is, the effect of small changes in the parameters. While these are extremely suggestive, they have the flaw that they depend on




Figure 3. Parameter hierarchies that span many orders of magnitude are observed in a wide variety of models. The vertical axis shows eigenvalues of the Fisher information matrix, scaled by the largest eigenvalue. Cell signaling data from [2], radioactive decay and neural network are taken from [53], quantum wavefunction are taken from [54], diffusion model and Ising model are taken from [38], meat oxidation is from [55], CMB data from [41], accelerator model taken from [56], van der Pol oscillator taken from [57], circadian clock model from [7], SW interatomic potential model from [58], power systems from [59], gravitational model from [60], transmission loss in an underwater environment from [61], and finally H_2IO_2 combustion model from [62].

the particular parameterization chosen. If we re-scale some parameters, such as by changing their units, then the FIM and hence the spectrum of eigenvalues will change⁸.

We can avoid this by measuring the global geometry of the model manifold [63, 64] using *geodesics* along the thick and thin directions. Let $\theta(\tau)$ be a path in parameter space, which then is mapped by the model onto a path on the model manifold. The length of this path, as defined by the Fisher information metric:

$$L = \int \sqrt{ds^2} = \int_0^1 d\tau \sqrt{\sum_{\mu\nu} g_{\mu\nu}(\theta(\tau)) \frac{\partial\theta^\mu}{\partial\tau} \frac{\partial\theta^\nu}{\partial\tau}}, \quad (6)$$

is independent of the way we define the coordinate system L on the model manifold. A geodesic is a curve connecting two points which minimizes the length between them. For models with Gaussian noise $x = y + \mathcal{N}(0, \sigma^2)$ as in equation (4) above, this geodesic distance agrees with the Euclidean path length in the prediction space, measured by walking along \mathcal{Y} .

We measure the widths of the model manifold by starting from an initial point (say, the best fit to the experimental data) and launching geodesics along one of the eigendirections in parameter space, shooting both ways until it hits the boundaries. Figure 4 shows the resulting geodesic widths across the model manifold for a number of different models. As advertised, they too show a roughly geometric progression over many orders of magnitude, with the stiff parameter directions

yielding the widest directions and with the many sloppy directions corresponding to incredibly thin directions on the model manifold. We will call this emergent hierarchical structure of the model manifold a *hyperribbon*, after the idea that a ribbon is much longer than it is wide, and much wider than it is thick.

We call a model sloppy, and the model manifold a hyperribbon, if the parameter-space eigenvalues and prediction-space widths are hierarchical: roughly equally spaced in log, spanning several decades. A sloppy model will have an effective low-dimensional description if many of its widths are small compared to the precision of the predictions demanded. This need not be the case—the Λ CDM model of the Universe cosmic microwave background radiation (column 3 in figure 3 and column 6 in figure 4) is sloppy (especially considering how few parameters it has) because the experiments are so precise that all parameters are well determined from the data. An impressive use of optimal experimental design [65] has shown that the individual parameters in our cell signaling network [1, 2] (column 1 in figure 3) could in principle be determined, albeit with a rather large number of experiments [66]. Alternatively, one could measure all the individual parameters separately in a model whose *collective* behavior is sloppy [67], but the resulting collective behavior will usually be predictable only if every one of the parameters is well specified [68]. Finally, most scientific models of complex systems like biological signaling are likely approximations to the microscopic behavior (Michaelis–Menten equations replacing enzyme reaction kinetics), and the resulting parameters are ‘renormalized’ to incorporate coarse-grained ignored complexity. Systematically designing experiments to measure these coarse-grained model parameters may just uncover flaws in the microscopic description that are irrelevant to the

⁸ In fact, the process of diagonalizing the FIM to find its eigenvectors amounts to finding a basis for parameter space, near a point, in which the matrix becomes 1.




Figure 4. Manifold widths for many disparate, nonlinear models (rescaled by the largest width for each model) illustrating the hyperribbon structure that characterizes model manifolds. Note the enormous range in vertical axis. For probabilistic models, imaginary lengths (i.e. negative squared distances) are reflected by dashed lines (section 7). Exponential curves, reaction velocities, and epidemiology model taken from [30], biased coin, Gaussians, CMB, Ising Model and Neural Network taken from [41]. The first three models here, and models 1–3 and 7–17 of figure 3, are nonlinear least-squares models, hence subject to the rigorous hyperribbon bounds of section 3. The other probabilistic models show the same hyperribbon hierarchy; methods for visualizing their model manifolds are discussed in section 7. Section 7.3 provides a different hyperribbon embedding that yields a finite-dimensional embedding for all of the latter except CMB (zero widths beyond the first few).

emergent behavior. We shall systematically study these useful coarse-grained models in section 4.1.

3. Hyperribbon bounds for nonlinear least-squares models

A fundamental feature of sloppy models is the hyperribbon structure of the model manifold. As discussed in the previous section, the model manifold has a hierarchy of widths that follows a geometric spread—roughly even in log. Where does this ubiquitous hierarchy of sensitivity come from? In this section we shall review rigorous bounds on the manifold of possible predictions for the special case of a multiparameter nonlinear least-squares model [30]. The proof relies on the smoothness of the predictions as the experimental controls (time, experimental conditions, etc) are varied, and explains the hierarchy of widths of the model manifold (as in figure 4). The smoothness and analyticity of the functions used in model construction directly controls the range of possible predictions allowed by the model. By quantifying the former we can understand the hyperribbon nature of the latter.

Before we start, we should emphasize how the sloppiness we study is distinct from other well-studied mechanisms that can make systems ill-posed. First, in systems biology and other fields, many have studied *structural identifiability*—when certain parameters cannot be gleaned from a given set of experiments, e.g. when some perhaps nonlinear transformation of the parameters makes for exactly the same predictions [23]. These would give eigenvalues of the metric exactly equal to zero. Our

work instead focuses on what one would call *practical identifiability* [23], where some parameter combinations are almost interchangeable, leading to the small, sloppy eigenvalues in parameter space. Second, while the models we study often have a variety of time or length scales, the observed hierarchies are *not* primarily due to a separation of scales [69]. Systems with a separation of scales (fast-slow dynamical systems, singular perturbation theory, boundary layer theory) have been thoroughly studied in the literature, and certainly will generate a large range of eigenvalues, but do not explain why eigenvalues span many decades uniformly. To test that separation of scales is not required for sloppiness, one need simply take a classically sloppy model such as a complex network in systems biology [1, 2] and set all the rate constants to one. The cost Hessian remains sloppy, with eigenvalues roughly uniformly spread over many decades. We also note that models with diverging time scales will typically violate the smoothness assumptions we need to derive our bounds. That is, a separation of time scales *adds to* the hierarchical sloppiness we derive in this section.

3.1. Taylor series

We can understand sloppiness, and give bounds for the widths of the model manifold that explain the observed hierarchies, through approximation theory. Specifically, we study polynomial approximations of model predictions. For simplicity, here we consider models whose predictions are a one-dimensional function $y_{\theta}(t)$ over the interval $t \in (-T, T)$, such as in figure 1

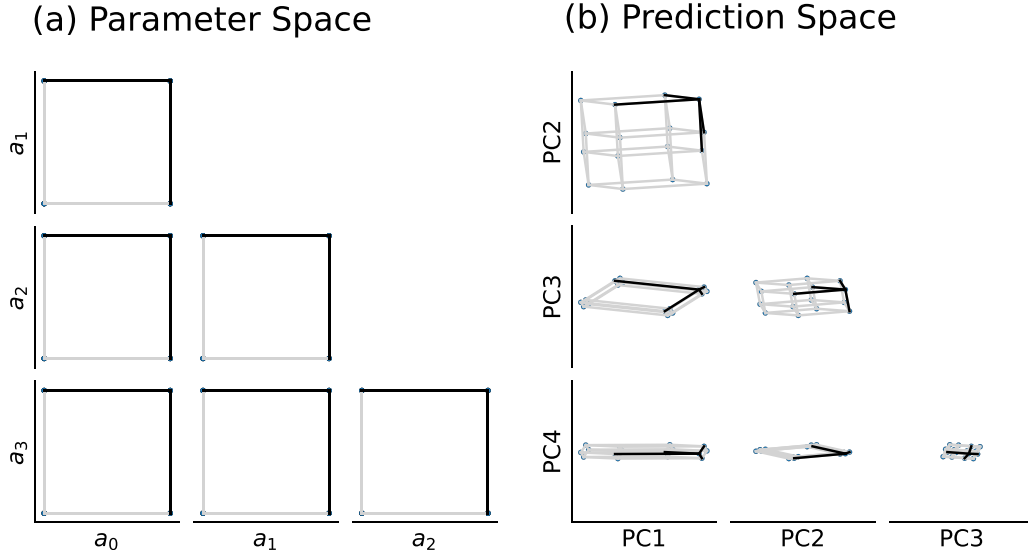


Figure 5. Model manifold for \mathcal{F} , projected along various principal axes of V . (a) The coefficients of $f_{N-1}(t)$ are constrained such that $|a_i| \leq 1$. They therefore form a region in parameter space for \mathcal{F} in the shape of a hypercube. (b) When evaluated at points $\mathbf{t} = [-0.9, -0.5, 0, 0.4, 0.8]$, the set of predictions are constrained by a hyper-parallelepiped whose widths are determined by V in equation (8), which we visualize by projecting along the Principal Components. Because every point in the model manifold y_θ is at most ϵ_N away from a corresponding f_{N-1} , the shape in (b) constrains the model manifold \mathcal{Y} , forcing a flat hyper-ribbon structure. Black lines show how the right angles in a vertex of the hypercube become highly skewed in prediction space.

where we choose $T = 0.9$. For models where this is not the case, we can simply rescale t so that it fits in this region. The bounds on model predictions shall rely on the smoothness of y as the experimental control t is varied.

In the simplest case, assume y_θ is analytic in t , with Taylor series $\sum a_k t^k$ and coefficients $|a_k| \leq 1$ for all k . This is a sufficient condition to give the series radius of convergence equal to one, larger than the radius T of the time range being studied. Consider approximating $y(t)$ by truncating its Taylor series. The approximating polynomial, $f_{N-1}(t) = a_0 + a_1 t + \dots + a_{N-1} t^{N-1}$, differs from $y(t)$ by at most:

$$\epsilon_N = \sum_{N}^{\infty} |a_k| t^k \leq \sum_{N}^{\infty} T^k \leq T^N / (1 - T), \quad (7)$$

since $|a_k| \leq 1$ and $|t| < T$. Recall from equation (1) that the model manifold \mathcal{Y} is defined as the space of all possible model predictions y_θ for varying model parameters θ . Understanding the constraints on f_{N-1} for each y_θ will therefore allow us to understand what constraints are set on the model manifold \mathcal{Y} . Due to the construction of f_{N-1} , we have the following linear equation:

$$\begin{aligned} \begin{bmatrix} f_{N-1}(t_0) \\ f_{N-1}(t_1) \\ \dots \\ f_{N-1}(t_{N-1}) \end{bmatrix} &= \begin{bmatrix} a_0 + a_1 t_0 + \dots + a_{N-1} t_0^{N-1} \\ a_0 + a_1 t_1 + \dots + a_{N-1} t_1^{N-1} \\ \dots \\ a_0 + a_1 t_{N-1} + \dots + a_{N-1} t_{N-1}^{N-1} \end{bmatrix} \\ &= \begin{bmatrix} 1 & t_0 & \dots & t_0^{N-1} \\ 1 & t_1 & \dots & t_1^{N-1} \\ \dots & \dots & \dots & \dots \\ 1 & t_{N-1} & \dots & t_{N-1}^{N-1} \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ \dots \\ a_{N-1} \end{bmatrix}, \end{aligned}$$

or, in particular:

$$f_{N-1}(\mathbf{t}) = V(\mathbf{t})\mathbf{a}, \quad (8)$$

where V is the well-studied Vandermonde matrix $V_{jk} = (t_j)^k$. The best-known property of the Vandermonde matrix is its determinant $\det(V) = \prod_{0 \leq i < j < N} (t_i - t_j)$. Notice that, since $t \in (-T, T)$, there are $\sim N^2/2$ factors $|t_i - t_j|$ ranging in size from $2T$ down to $1/N$: the determinant (the product of the eigenvalues) will be very small.

Because the model predictions are characterized by the coefficients in their Taylor series as well as where they are evaluated (i.e. the a_i and $V(\mathbf{t})$ in equation (8) respectively), understanding what constrains both a_i and $V(t)$ will provide an understanding of what bounds the model.

For instance, for our models where the a_i are limited to be between -1 and 1 , the Vandermonde matrix takes the hypercube of possible values of \mathbf{a} and maps it to a hyper-parallelepiped \mathcal{F} in prediction space that is incredibly skewed, as shown in figure 5; its volume is proportional to the incredibly small determinant. The model manifold \mathcal{Y} is constrained to be within $\epsilon_N = T^N / (1 - T)$ (equation (7)) of this hyper-parallelepiped because each $y_\theta \in \mathcal{Y}$ has a truncated power series f_{N-1} which differs from it by at most ϵ_N . The highly skewed nature of this hyper-parallelepiped, whose widths are determined by the Vandermonde matrix, forces a similar hierarchical structure on the model manifold \mathcal{Y} .

3.2. Chebyshev polynomials

To expand the understanding of why a hyperribbon structure is observed for non-analytic models, in [30] it is proven for a relaxed smoothness condition where the coefficients a_i are

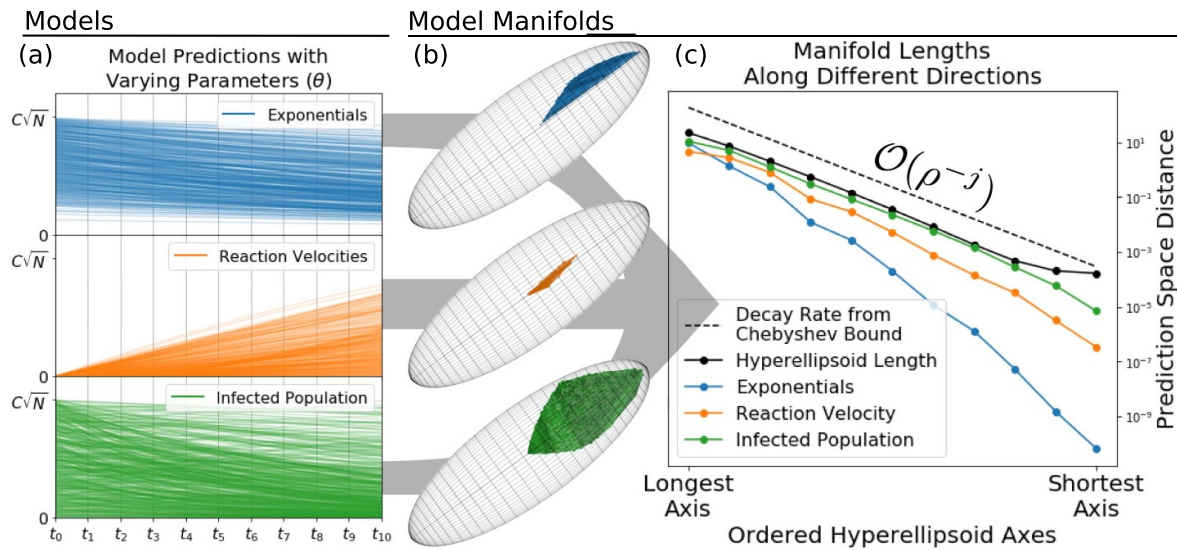


Figure 6. Model manifold and bounds of three disparate models: (1) exponential curves such as those in figure 1, (2) reaction velocities of an enzyme-catalyzed reaction, and (3) the infected population in an SIR model. (a) Each model is evaluated at specific points t_i for varying parameter values, yielding a set of different predictions. (b) Each model manifold is bounded by a the same hyperellipsoid, whose widths follow a geometric decay ($O(\rho^{-j})$), calculated from the singular values of the Vandermonde matrix combined with an error from Chebyshev Approximation. (c) The widths of the hyperellipsoid are plotted against the successive widths of each model manifold, showing a clear hyperribbon structure. The smoothness of the models determines the value of ρ , which in turn sets the widths of the hyperellipsoid bound, which forces a hyperribbon structure for the model manifolds. Reprinted figure with permission from [30], Copyright (2019) by the American Physical Society.

bounded not by a hypercube but by an n -sphere of radius r , and not for all coefficients but only those up to a certain order. Bounds were derived both using a Chebyshev approximation and by using the Taylor series considered above. Following a very similar line of reasoning, it is possible to show that the model manifold \mathcal{Y} is constrained by a hyper-ellipsoid whose j^{th} largest principal axis length is bounded by a power law ρ^{-j} . To illustrate this, consider three nonlinear models: (a) a sum of exponentials, such as in figure 1, (b) reaction velocities of an enzyme-catalyzed reaction, and (c) the infected population in an SIR epidemiology model. In all three cases, the models have varying model parameters which produce different predictions, shown in figure 6(a). The model manifolds are all bound by the same hyperellipsoid, shown in figure 6(b), and form a hyperribbon structure as the successive widths of the hyperellipsoid follow a geometric decay.

Hyperribbons have also been observed for models with multiple control variables, for instance those which vary for both time and temperature [30, supplemental material]. Here, a two-dimensional Vandermonde matrix is constructed to understand the bounds in prediction space, producing a similar decay in the widths of the hyperellipsoid that bounds the model manifold with additional degeneracies based on the number of additional control variables.

3.3. Interpolation theory

We can get a different intuitive picture for why smoothness in the control parameters implies a hyperribbon structure in the predictions, by considering interpolation theory [63, 64]. Consider studying successive cross sections of the model

manifold generated by experimental measurements at times t_0, t_1, \dots , here for analytic functions whose radius of convergence $|y_\theta^{(n)}(t)|/n! < 1$ for all points t in the interval [53, 63]. The value of $y_\theta(t)$ at each time $t \in (-T, T)$ is an experimental prediction, and therefore a coordinate axis of our model manifold. Because there are an infinite number of points on the interval, and each point lends itself to a dimension on the model manifold, the resulting manifold is infinite-dimensional. Because of the bounded derivatives⁹, measuring $y_\theta(t)$ at a new time further constrains the range of possible values at other times, as shown in figure 7. Each experimental measurement $y_\theta(t_j)$ fixes a point, and has the effect of slicing the model manifold along axis t_j at y_j —keeping only the parameters θ whose predictions agree with the experiments.

Figure 7 shows a function $y_\theta(t) \in \mathcal{Y}$ taking values $y_\theta(t_j) = y_j$, and rigorous bounds on any function $z(t) \in \mathcal{Y}$ that agree with $y_\theta(t)$ at the constrained time points. The range of predictions $z(t')$ in figure 7 gives the width of the model manifold cross section after slicing along certain fixed points. Note

⁹ The function $y_\theta(t)$ in figure 7 was chosen by generating random quintic polynomials $\sum_{\mu=0}^5 \theta_\mu(t)^\mu$ with bounded monomial coefficients $|\theta_\mu| < 1$, and then selecting for ones whose Taylor expansions at all other points in the interval have the same bound (hence $\max_{t,n} |y_\theta^{(n)}(t)|/n! < 1$). This further selection allows for application of the standard interpolation theorem [63]. Both conditions correspond to forcing a radius of convergence equal to one, slightly larger than the radius of our interval $t \in (-0.9, 0.9)$. But the second criterion, forcing a bounded derivative, leads to a constraint in the allowed space of subsequent predictions which in turn creates the hyperribbon structure.

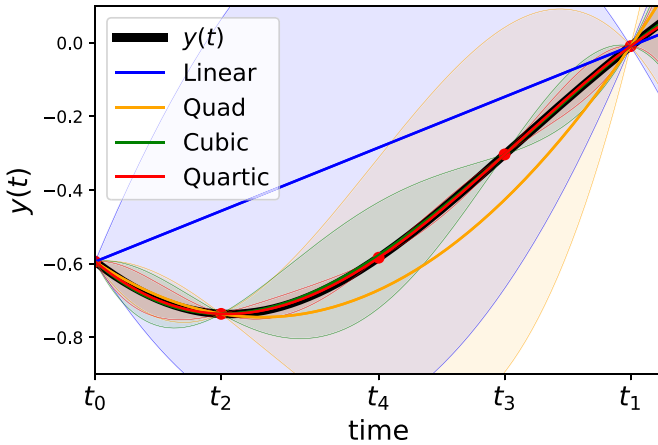


Figure 7. Convergence of interpolations. The black curve is a model prediction for a one-dimensional function $y_\theta(t)$ for $-0.9 \leq t \leq 0.9$. The colored curves are polynomial fits $f_{n-1}(t)$ as we add data points $y_\theta(t_n)$ at new interpolating points t_n . Note that the widths of possible errors in the fits (shaded regions) get smaller as we add more points: the interpolations converge with more data. The shaded regions give the usual error bound on the interpolation, given the smoothness of $y_\theta(t)$. Our function has $y^{(n)}(t)/n! < 1$ with radius of convergence at any point slightly larger than our interval. The geometrical decrease in the widths of these bounds with n is related to our hyperribbon bounds on the widths of successive cross sections of the model manifold.

that this width shrinks with each new experimental measurement, as the effect of additional constraints only further limits the range of allowed functions. Fixing the value of $y_\theta(t_j) = y_j$ shrinks the range of all other possible predictions, just as a ribbon snipped in two has a cut edge that typically reflects the width and thickness of the original ribbon. The convergence of a polynomial interpolation between the points $y(t_j)$ is equivalent to a hyperribbon structure of the model manifold \mathcal{V} .

4. Effective theories as manifold boundaries

A remarkable insight that comes out of the geometric study of multi-parameter models is how models of varying complexity relate to one another. Simple, low-dimensional models often reside on the boundaries of complex, high-dimensional ones. We have already alluded to this idea for the motivating example in equation (2). To illustrate constructions that will be central to our approach, consider the hierarchy theories to emerge after the discoveries of relativity and quantum mechanics.

To begin, consider a multi-parameter model formulated in the context of special relativity and include the speed of light, c , among the parameters. Varying c defines a sequence of models of co-dimension one. In the limit of this sequence $c \rightarrow \infty$ the predictions of the model reduce to those of classical Newtonian spacetime. This is, of course, one of the well-known classical limits and formalizes the relationship between relativity and classical kinematics. In geometric language, classical kinematics is a boundary of the manifold of special relativities. Similarly, the limit of $\hbar \rightarrow 0$ identifies classical mechanics as a boundary of quantum mechanics. Including G

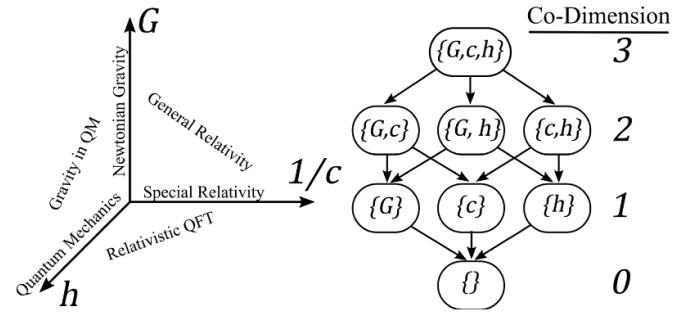


Figure 8. The physical constants G , c , and \hbar parameterize a hierarchy of effective theories connected by a directed graph of limiting approximations.

as a parameter completes the three-dimensional space of fundamental theories sometimes known as cGh physics [70].

Similarly, manifolds of co-dimension two correspond to models built in the context of general relativity ($\{G, c\}$), relativistic quantum field theory ($\{c, \hbar\}$), and (less well-known) classical gravity in quantum mechanics ($\{G, \hbar\}$) [71, 72]. These manifolds are bounded by the simpler theories that preceded them and presumably are themselves limiting cases of a more fundamental theory parameterized by all three. This topological relationship among theories is summarized by the directed graph on the right in figure 8.

These connection between theories are well-known to physicists and the geometric constructions presented here is not new [71–73]. Indeed, this relational interpretation of effective theories is a key in formalizing the criteria under which a model is valid and which are likely to be the leading order corrections. Einstein and Schrödinger tell us when and why Newton is right. Furthermore, this way of organizing scientific knowledge is useful for many types of multi-parameter models. For example, among continuum models of conventional superconductivity, the detailed theory of Eliashberg involve many greens functions detailing coupling information between electrons and phonons [74]. From this, a hierarchy of simplified models emerge as limiting cases, i.e. boundaries of a more complicated model manifold. The weak-coupling limit gives the theory of Gor’kov [75], the semi-classical limit gives that of Eilenberger [76], and the high-temperature limit gives Ginzburg–Landau theory [77].

Other domains have similar hierarchies of theories and models. We find this structure generic in multiparameter models. It motivates an algorithm for systematically constructing effective theories: the Manifold Boundary Approximation Method (MBAM).

4.1. The manifold boundary approximation method

Simplified models are found on the boundaries of the model manifold, and these boundaries exhibit a universal structure. Returning to the motivating example of two exponentials introduced in equation (2) and figure 1, observe that the boundary has a triangle-like structure: a two-dimensional manifold with three one-dimensional edges that meet at three zero-dimensional cusps. Although the size and shape of the model




Figure 9. The boundary complex of the model manifold is summarized by a directed graph called a Hasse diagram. Elements of the boundary complex correspond to simple, reduced models that are related by physically interpretable limiting approximations.

manifold depend on the observation time, this triangle-like boundary structure is independent of these details and depends only on the form of the observation function, $y_\theta(t)$. We have noted that each of these boundary elements is a simplified, approximate model. Formally, this boundary structure is a hierarchical cell complex that is a natural topological decomposition of the boundary.

Similar boundary complexes are found in models from very diverse contexts [78, 79]. That is to say, a typical N -dimensional model is bounded by a collection of $N - 1$ -dimensional faces, which join at $N - 2$ -dimensional edges, and so forth. The adjacency relationship among these facets of different dimensions defines a partially ordered set (often called a POSet), described as graded because each facet has a dimension associated with it. These relations can be conveniently visualized as a directed graph known as a Hasse diagram, as we have done in figure 8 for *cGh* theories and repeated in figure 9 for the exponential example.

Each of the boundary cells of a model manifold corresponds to a simplifying approximation that corresponds to an extreme, idealized behavior of the model. The Hasse diagram summarizes how these approximations can be composed to yield progressively simpler models. That is to say, the Hasse diagram of the boundary complex is a road map from a complicated multi-parameter model through various approximations to the distinct behavioral regimes the model enables. When the model manifold is endowed with an information-theoretic metric, it, along with its boundary structure, is compressed in sloppy directions and stretched in stiff ones. Useful approximations, i.e. those that retain key explanatory variables, correspond to those boundary cells aligned with the long-axes of the manifold.

We translate this observation into an algorithm, the *Manifold Boundary Approximation Method* (MBAM) that uses information geometry to identify simplified models as boundary cells of a model manifold that is appropriate for a particular data set. The basic idea is to construct a geodesic shortest-path through the model's parameter space until it intersects a face on the boundary complex. By choosing a sloppy direction in parameter space (searching both directions along an

eigenvector of the FIM with small eigenvalue), the geodesic will find a nearby boundary point with nearly identical predictions. By considering parameter values along the geodesic path, one deduces the limiting approximation that corresponds to the appropriate boundary cell. Often, this limiting approximation may involve one or more parameters going to extreme values, such as infinite or zero, in a coordinated way. We re-parameterize the model so that participating parameters are grouped together into a single combination that goes to zero at the boundary of the manifold. Finally, we analytically apply the approximation to the model to construct a simpler model with one fewer parameter. We iterate this process until the model is sufficiently simple.

MBAM has a formal similarity to using scale separation to perform model reduction. Scale separation is one of the oldest and most well-established methods of model reduction. It introduces a small parameter in which to perturb; the model reduces to a simplified form in the limit that the parameter becomes zero. Similarly, MBAM determines a combination of parameters, the perturbation parameter, that goes to zero at the boundary of the model manifold. However, sloppiness adds a new wrinkle to how we use and interpret this procedure. We do not necessarily require the true value of the perturbation parameter to be small. Rather, the true parameter value may be large provided the data to be explained do not constrain its value. Indeed, we observed above that although a model with separated scales will be sloppy, separated-scales are not necessary for sloppiness. The full model is statistically indistinguishable from the simpler separation-of-scales model. We advocate for using simple models because they reflect the informativity of the data, not because the true system to be modeled has well-separated scales.

4.2. Relation of MBAM to other reduction methods

Because the problems related to model complexity transcend scientific disciplines, there are many well-established reduction techniques that have grown-up around specific application areas. Most of these exploit the specific structure of the problem at hand, usually in conjunction with expert intuition. For example, dynamical systems are often simplified using singular perturbation theory in which the fast dynamics (often identified by an expert) are replaced by the slow manifold. Since, MBAM is agnostic to the mathematical form of the model, it can be compared to domain-specific reduction methods. In many cases, MBAM is formally equivalent to these techniques. However, because it is a data-driven method, it need not require expert insight and can guide insights into the modeling of complex systems.

For example, consider the problem of modeling transient dynamics of a power system. Power systems are large, engineered systems with many interconnected components. Time scales vary from speed-of-light electrical responses to load variations over hours and days. Models incorporate multi-physics in the form of heterogeneous components that describe not just power delivery, but generation from multiple sources, including renewables, and interactions with economic demand.

Of particular interest are models of synchronous generators, a power systems component that plays an important role in large simulations. The power community has developed a hierarchy of models describing synchronous generators that are formally connected through a series of singular perturbations. A typical inference problem involves estimating machine parameters from short disturbances that may last a few seconds. This problem is challenging because the relevant time scales are intermediate between the fastest and slowest time scales of the system, so it is difficult to judge which time scales are relevant. Indeed, the true dynamics are not necessarily faster or slower than the observation time scales, but parameters are unidentifiable because of sloppiness.

Applying MBAM to slowly-developing disturbances leads to the expected sequence of singularly perturbed reduced models advocated by theorists. However, when applied to disturbances at intermediate time scales, the results are more interesting. Recall that singularly perturbed systems have both an inner and an outer solution corresponding to the fast and slow dynamic respectively. Rather than uniformly removing fast dynamics, MBAM selects a combination of inner and outer solutions of the singularly perturbed system to remove parameters not constrained by the data. The remaining parameters are those that could be learned from the data.

Next, we turn to the theory of linear control that has produced many formal insights into the approximation of dynamical systems. We first discuss how control-theoretic quantities relate to sloppiness, and then demonstrate how MBAM relates to another standard approximation method known as *balanced truncation*.

A linear, time invariant (LTI) system takes the form:

$$\begin{aligned}\dot{x} &= Ax + Bu \\ y &= Cx + Du,\end{aligned}\tag{9}$$

where $x(t) \in \mathbb{R}^n$ are state space dynamical variables, $u(t) \in \mathbb{R}^p$ are inputs/control variables, and $y(t) \in \mathbb{R}^m$ are observed variables. A , B , C , and D are constant matrices of the appropriate shape. This model arises as the linearization of a generic dynamical system near a stable equilibrium, and theorems proved on LTI systems generalize to nonlinear dynamical systems under certain regularity conditions. Two important concepts to emerge in the study of LTI systems are observability and controllability. Roughly, observability is how well the internal state of the system can be inferred from its outputs while controllability is how easily a system can be driven to a particular state through a control protocol. The concepts of observability and controllability generalize to nonlinear systems, though they are most easily quantified in the context of LTI systems.

The model reduction problem as formulated in controls is not immediately analogous to the parameter reduction problem addressed by MBAM. For controls, ‘model reduction’ refers to reducing the dynamic order of a system. The goal is to approximate a system in the form of equation (9), with another LTI system with a smaller state space (i.e. $\tilde{x}(t) \in \mathbb{R}^{n'}$ where $n' < n$).

Balanced truncation uses the concepts of observability and controllability to identify accurate approximations. For LTI systems, observability and controllability are each quantified by a Gramian matrix, W_O and W_C respectively. The Gramians are symmetric, positive definite matrices that act as metrics quantifying the observability and controllability of vectors in the state space. A fundamental theorem for LTI systems shows that for an appropriate change of basis in the state space, the resulting observability and controllability Gramians are equal and diagonal. In this so-called ‘balanced representation,’ the vectors in the state space are equally controllable and observable. *Balanced truncation* is a reduction strategy in which an LTI system is first transformed into its balanced representation and the least observable/controllable states are truncated (removed) from the system.

To relate state-space reduction methods of control theory to MBAM, we parameterize the family of LTI systems of fixed dynamical order n . We generated data by choosing many potential inputs $u(t)$ and evaluating the resulting time series $y(t)$. The resulting model manifold is sloppy, and the sloppy *parameters* are those components in the A , B , and C matrices that determine the dynamics of the least observable/controlable *states*. When we apply MBAM to this system, we find the manifold boundary corresponding to the limit that the unobservable and uncontrollable states are constant decouple from the observed variables, $y(t)$. This limit is equivalent to the balanced truncation method advocated by the controls community.

Balanced truncation has two major limitations. First, it is difficult to interpret because physically meaningful states are mixed together in the balanced representation. It is also difficult to generalize to nonlinear systems. However, the information geometry techniques underlying sloppy model analysis naturally accommodate these regimes. Thus, MBAM is a natural generalization of principles of balanced truncation to structure-preserving and nonlinear model reduction of interest to the control community.

4.3. Beyond manifold boundaries

As presented here, MBAM has some limitations and raises some fundamental questions about the geometric and topological properties of multi-parameter models. Here we discuss the relationship between MBAM and more general statistical theory. We also highlight some interesting formal considerations related to the parameter singularities that are beyond the scope of this review.

In statistics, simplified models that are special cases of a more general model are called nested models. Nested models are important in many statistical applications, especially hypothesis testing. The boundary models that we describe here are a specific type of nested models. But more generally, sometimes useful reduced models may also be found at non extreme values. For example, in a spin system with parameters quantifying couplings (e.g. nearest neighbor, next nearest neighbor, etc) that can be either positive or negative, a natural reduced model is to fix these couplings to zero rather than to extreme values. Future work could adapt the MBAM to identify other

types of useful reduced models that are not on the boundary. Whether there are geometric ways to identify these other nested models that do not make explicit reference to their having one less parameter remains an open question.

Taking parameters to extreme values is a common reduction strategy in science because the reduced model can often be evaluated in a different way. For example, in a singularly perturbed dynamical system, an ordinary differential equation transitions into a differential-algebraic equation [80]. In general, the singular nature of these limits makes them especially relevant for inter-theoretic reasoning in mathematical physics as in figure 8. More broadly, however, model manifolds may have singular models that are not strictly boundaries, such as at the tip of a cone. These singularities often correspond to useful reduced models. More work is needed to understand the geometric and topological properties of these generalized singularities and develop reduction algorithms to efficiently find them.

5. Bayesian priors which perform well for sloppy models

A parameterized model $p(x|\theta)$ represents predictions for the likelihood of seeing data x in some experiment, given parameter values θ . The simplest question in Bayesian analysis is the reverse: after seeing data x , can we infer the parameters? Bayes' rule gives us not a single answer, but a distribution, proportional to a *prior distribution* $p(\theta)$:

$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{p(x)},$$

where the denominator is simply a normalization constant. This result is just mathematics, but deciding how to choose a prior (and exactly what this means) has long been a delicate issue [81]. In some contexts we may have detailed pre-existing knowledge, and if we express this as a relatively narrow distribution $p(\theta)$, then Bayes' rule represents a small update. But when we know little, we should express our ignorance as a broad distribution, some kind of *uninformative prior*. What concerns should guide us in formalizing this idea, and how much do they matter?

One requirement is that put forward by Jeffreys: all results should be reparameterization invariant, and hence priors must transform like a density. For example, they should not depend on whether we choose to parameterize a model of exponential decay (like equation (2)) with decay rates or with their inverses, lifetimes. This would not be true for a flat prior (meaning equal weight in parameter space), but it is true for Jeffreys prior, which is simply the volume form arising from the Fisher metric, up to normalization [82].¹⁰

¹⁰ Despite bearing his name, this was not what he favored using in all situations. He argued for special treatment of location parameters, meaning directions along which the entire system is exactly symmetric, omitting them from the dimensions over which this volume form is calculated. The models considered here have no such parameters.

$$p_J(\theta) \propto \sqrt{\det g_{\mu\nu}(\theta)}. \quad (10)$$

This volume measure encodes a continuum view of the manifold, in which it is infinitely divisible. There is no notion that some parameter directions could be too short (too sloppy) to measure: Scaling the metric $g_{\mu\nu} \rightarrow M g_{\mu\nu}$ has no effect at all on $p_J(\theta)$ since M is absorbed by normalization. This scaling of distances by \sqrt{M} is precisely the effect of repeating our experiment M times, $p(x|\theta) = \prod_{i=1}^M p(x_i|\theta)$, and with enough repetitions, we can observe effects previously lost in the noise. This continuum view makes $p_J(\theta)$ well-adapted for models in the limit of infinitely much data, but as we argued above, the multiparameter models appearing in science are very far from this limit. Here, we argue that using it for such models can lead to extreme bias in relevant results. Geometrically, the problem is that it places all of its weight in regions where the hyperribbon is thickest in sloppy and thus unmeasurable directions.

A better behavior on a hyper-ribbon would be to place weight only according to thickness along measurable, stiff directions, perhaps by being non-zero only on appropriate edges of the manifold (section 4.1). A prior that did this, schematically $p(\theta) = p_{\parallel}(\theta) \prod_{\mu \in \perp} \delta(\theta_{\mu})$, would encode a lower-dimensional effective model. This section describes some priors with this preference for edges. They open a view of model selection as part of prior selection.

5.1. Maximizing mutual information

The natural measure of how much information we learn from an experiment is the average reduction in the entropy over parameters θ upon seeing data x , which is the mutual information [83]:

$$I(X; \Theta) = S(\Theta) - S(\Theta|X) = \sum_x \int d\theta p(x|\theta) p(\theta) \log \frac{p(x|\theta)}{p(x)}.$$

This information depends on the prior. Maximizing it may be thought of as choosing a minimally informative prior, and this choice is one way to formalize what the notion of 'uninformative' means¹¹. The idea of choosing a prior by maximizing this information was studied by Bernardo, but always in the limit of $M \rightarrow \infty$, infinitely many repetitions or infinitely much data, where the resulting priors are called reference priors [84]. In this limit they approach Jeffreys prior (under mild conditions), and this idea may be seen as an alternative motivation for using Jeffreys prior [85]. But away from this limit, the resulting priors behave very differently, without losing reparameterization invariance. And, in particular, they have model selection properties which have been overlooked in the statistics literature.

Priors which maximize mutual information were investigated in [36], which showed that they have a preference for edges, and hence tend to select simpler subspaces. In fact they

¹¹ Minimizing instead of maximizing, notice that a prior with all of its weight at any one point has mutual information $I(X; \Theta) = 0$, the lowest possible value. Such a prior represents complete certainty, and it seems fair to call it maximally biased.

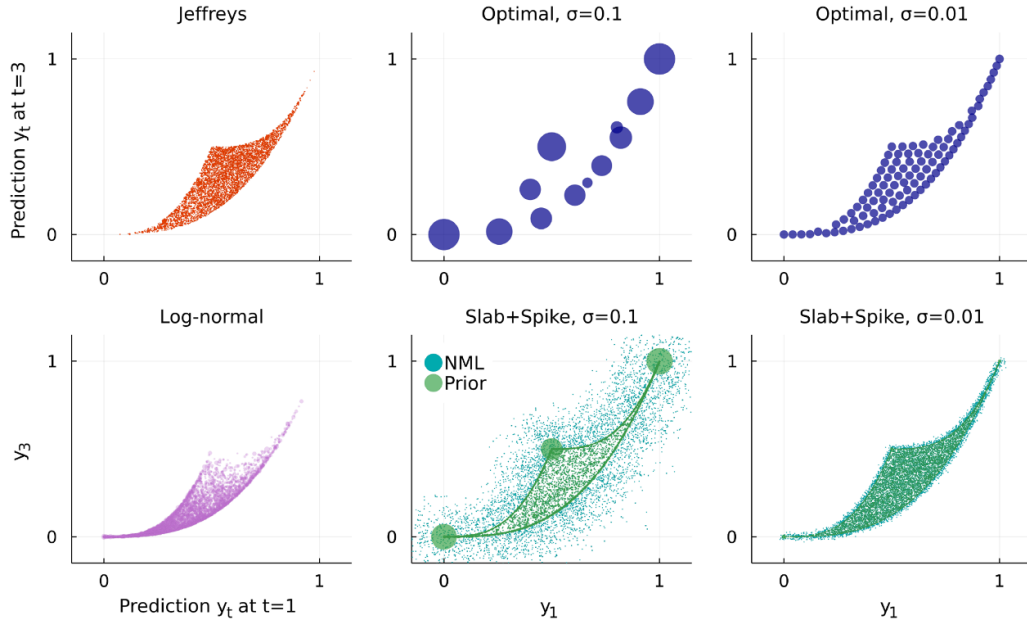


Figure 10. Various priors for $D = 2$ sum-of-exponentials models. While Jeffreys prior (top left) is independent of σ , the discrete optimal prior $p_*(\theta)$ (upper middle and right) has weight on zero- and one-dimensional edges at large noise, but adjusts to have weight in the two-dimensional interior as the noise is decreased. The plots use as co-ordinates $y_t(\theta)$ for $t = 1, 3$, which is an isometric embedding onto the paper. The adaptive slab-and-spike prior $p_\#(\theta)$ (lower right, discussed in section 5.2) also depends on the noise level. At large σ , $p_{\text{NML}}(x)$ has most of its weight outside the model manifold, hence $p_\#(\theta)$ has most of its weight on the edges (and uniform inside). At small σ , it approaches Jeffreys prior.

are almost always discrete, with weight on K atoms (points in parameter space) [86–89]:

$$p_*(\theta) = \underset{p(\theta)}{\operatorname{argmax}} I(X; \Theta) = \sum_{a=1}^K \lambda_a \delta(\theta - \theta_a). \quad (11)$$

Along a short dimension (say $\lesssim 1$ standard deviations) the atoms will lie only at the boundaries. But along a long dimension ($\gg 1$ standard deviations) there will be many atoms, approaching a continuum. This turns out to be precisely the behavior we want on a hyperribbon.

To illustrate this behavior, consider the sum-of-exponentials model equation (2), used in figure 1, except with more parameters (and normalized to one at $t = 0$):

$$x_t \sim y_t(\theta) + \mathcal{N}(0, \sigma^2), \quad y_t(\theta) = \sum_{\mu=1}^D \frac{1}{D} e^{-tk_\mu(\theta)}, \quad (12)$$

$$0 \leq k_\mu(\theta) \leq \infty.$$

Figure 10 shows $p_*(\theta)$ for a model with $D = 2$ parameters at various values of σ , drawn in y_t space for times $t = 1, 3$. (This choice of co-ordinates gives an isometric embedding: Euclidean distance on the paper is proportional to the Fisher metric.) At very large σ there are just two atoms at the extreme values (i.e the zero-dimensional boundaries). At very small σ , the data-rich regime, there are many points almost tiling the space. And in the limit $\sigma \rightarrow 0$ (equivalent to $M \rightarrow \infty$ repetitions) this approaches Jeffreys prior, which is flat in y . But

it is the intermediate values which represent the more typical hyper-ribbon situation, in which some dimensions are data-rich, and some are data-starved.

The idea of a prior which depends on the number of repetitions (or equivalently on the noise level σ) breaks the dictum ‘yesterday’s posterior is tomorrow’s prior’ for Bayesian updating. However, any prior which depends on the experiment to be performed (including Jeffreys) already breaks this rule. If tomorrow’s experiment is upgraded to have less noise in some components, even if it otherwise explores the same model, then the combination requires a new prior, thus the update from today to tomorrow is different from a simple application of Bayes rule. For Jeffreys prior there is an exception for repetitions of the same experiment.

But, as we demonstrate below, this invariance to repetition means that Jeffreys performs very badly for hyper-ribbon model manifolds, sloppy models, in which we cannot fix all parameters [37]. Indeed, the distinction between a parameter we can fix and one we cannot is precisely that we do not yet have enough data to fix the latter.

The idea of accepting a discrete prior is also initially surprising to classical intuition. What if the truth were between two points of weight, would we eventually have poor accuracy? Three replies are possible. First, ‘eventually’ here implies taking much more data, and further updating the posterior, which is ruled out by the prior’s dependence on the number of repetitions: the appropriate prior is the one for all the data. Second, along a relevant parameter direction, discretization is arguably no worse than truncating some measured θ to a

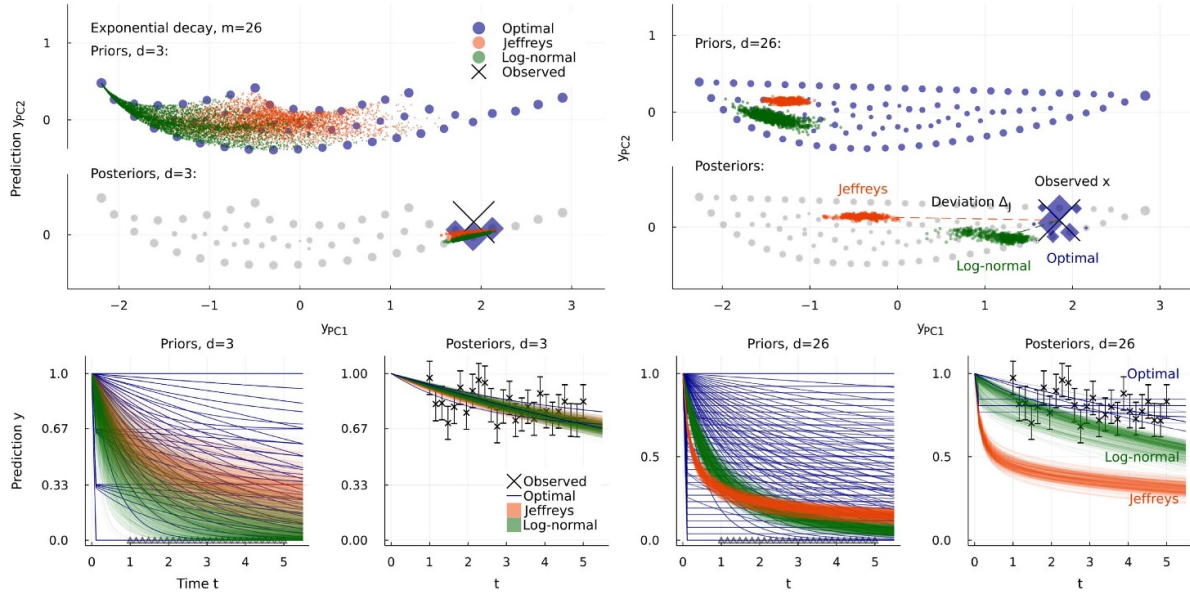


Figure 11. Comparison of maximum-information and Jeffreys priors, for $D = 3$ (left) and $D = 26$ (right) sum of exponentials models. Above, the priors in prediction space y , showing how Jeffreys concentrates weight in the center of the high-dimensional volume, while the points of the discrete optimal prior $p_*(\theta)$ remain spread out. Below, the priors and posteriors are drawn as time-courses y_t . The same observed data $x \in \mathbb{R}^{26}$ is used for both models; the error bars indicate the scale of the noise, which can be compared to the spacing of discrete lines. Note how badly Jeffreys prior distorts the posterior center of mass (upper right) and the inferred time-course (lower right) for $D = 26$. Reproduced with permission from [37].

sensible number of digits, for given error bars¹². And third, along an irrelevant direction, placing all weight at an extreme value is precisely what effective theories do. For example, QED assumes that the coupling constants of all forces except electromagnetism are exactly zero, which we know to be false in reality. But this false assumption is irrelevant to the experiments being described, and allows predictions of very high accuracy [91].

Jeffreys prior is, as we mentioned, the limit $M \rightarrow \infty$ of the maximum-information prior $p_*(\theta)$. So according to our arguments, it is the correct thing to use in the regime when all parameters are well-constrained, when all are in the data-rich regime. How badly wrong is it when we are not in this regime? The eigenvalues of the metric tensor in sloppy, multi-parameter models span an enormous range, so the square root of the determinant giving Jeffreys prior (equation (10)) is often enormously small. Worse, it varies often by many orders of magnitude even for small changes in parameters. Uninformative priors are supposed to be broadly agnostic about parameter choices, and Jeffreys prior fails in this regard.

Thanks to progress in numerical techniques for finding $p_*(\theta)$ [37, 92] we can now illustrate this extreme weighting. Figure 11 shows the same sum-of-exponentials model, but now comparing a model with $D = 3$ parameters to one with $D = 26$ parameters. Both describe the same data, measuring decays at $M = 26$ different times in $1 \leq t_i \leq 5$, each with noise $\sigma = 0.1$. The discrete priors for these two models are fairly

similar, but the concentration of weight in Jeffreys prior is much stronger in the higher-dimensional model. This is most clearly seen in the posterior, after observing a particular x as shown. In $D = 3$ Jeffreys performs well, but in $D = 26$ its posterior is centered 20 standard deviations away from the data. This is not a subtle error.

How many repetitions would be required in order for Jeffreys prior to be justified here? The smallest eigenvalue of the FIM is about 10^{-50} here. To make this order 1, we need more repetitions than there are seconds in the age of the Universe [37]. While this is obviously a toy model, tens of irrelevant parameters, and astronomical timescales to measure them all, are also found in more realistic models (figures 3 and 4).

Another way to compare these priors is by how much information we will learn. The scores (expressed in bits) for the priors in figure 11 are shown in figure 12. Notice that adding additional parameters almost always increases $p_*(\theta)$, but we observe that the extreme concentration of weight gives Jeffreys prior a much lower score in the high-dimensional model. Even without being able to find $p_*(\theta)$ exactly, this measure offers a tool for model selection. But unlike traditional approaches to model selection, this takes place before seeing data x . It depends only on the resolution of the experiment, not its outcome.

5.2. Adaptive slab-and-spike priors

Finding $p_*(\theta)$ is not always easy, as $I(X; \Theta)$ involves both an integral over X and a sum over the atoms: The K points of weight all interact with each other, via the data they generate.

¹² In fact the optimal prior is closer to continuous than this sentence suggests. The number of atoms K grows slightly faster than the length of a dimension L , by a novel power law [90].




Figure 12. Mutual information $I(X; \Theta)$ captured when using various priors for sum-of-exponentials model, against the number of model parameters D . Jeffreys and the optimal prior for $D = 3$ and $D = 26$ are drawn in figure 11 above, while the projected maximum likelihood prior is drawn in figure 13 below. Notice how close the projected prior is to the optimal prior in high dimensions, and how poorly Jeffrey’s and the log-normal priors perform by this measure.

We present here another prior which shares many of the virtues (and avoids Jeffreys’ vices), while being easier to draw samples from.

To derive this, consider first the following distribution on data space X , called a normalized maximum likelihood¹³:

$$p_{\text{NML}}(x) = \frac{\max_{\hat{\theta}} p(x|\hat{\theta})}{Z}, \quad Z = \int dx \max_{\hat{\theta}} p(x|\hat{\theta}). \quad (13)$$

For models with Gaussian noise, this describes a cloud of points filling and surrounding the embedded model manifold $\mathcal{Y} \subset \mathbb{R}^M$ (figures 10 and 13), with a range given by the measurement error in the predictions. We then define a prior distribution on parameter space Θ , by simply projecting the weight of $p_{\text{NML}}(x)$ at each x to the corresponding maximum likelihood point $\hat{\theta}$. We may write this as:

$$p_{\#}(\theta) = \int dx p_{\text{NML}}(x) \delta(\theta - \underset{\hat{\theta}}{\operatorname{argmax}} p(x|\hat{\theta})). \quad (14)$$

This prior has full support, but places extra weight on points on the boundaries of the model manifold, and more weight on lower-dimensional boundaries (such as fold lines, and corners), thus generalizing the idea of a slab-and-spike prior [95] to many dimensions and curved manifolds. How much weight is on boundaries is chosen automatically, depending on the level of noise in the likelihood function.

In practice, this prior can be computed efficiently by sampling $p_{\text{NML}}(x)$ using Monte Carlo methods, and then recording for each accepted point x the corresponding $\hat{\theta}$ to give a sampling distribution $p_{\#}(\theta)$ in parameter space. The Metropolis weight does not require the normalization factor Z .

¹³ This is also used, independently, in a non-Bayesian approach to model selection by minimum description length [93, 94]. In that context there is no measure on Θ , although a measure on X is needed to calculate the normalization Z . The normalization Z is not crucial, and indeed can be infinite (an *improper* prior); such a prior can still be sampled with Monte Carlo (see below).

Figures 10 and 13 show the result, in the same sum of exponentials models as above. For all models $x \sim y(\theta) + \mathcal{N}(0, \sigma^2)$, the embedding of the model manifold is smeared out by a distance σ to make $p_{\text{NML}}(x)$. This means that, unlike Jeffreys prior, the difference between places where the manifold is thin compared to σ , and places where it is very very thin, will not matter at all. The smeared weight is pulled back to contribute to $p_{\#}(\theta)$ exactly at the closest edge. And lower-dimensional edges have more weight; zero-dimensional points usually have finite weight. Thus $p_{\#}(\theta)$ is as a generalized slab-and-spike prior, with weight both in the volume and on edges.

Like $p_{\star}(\theta)$, the degree of preference for edges depends on the level of noise σ , or the number of repetitions. Along irrelevant, sloppy, short, directions, almost all the weight is on the edges. Along relevant, stiff, long, directions, most of it fills the interior. As the noise is decreased, dimensions will smoothly pass from the former category into the latter. We view this as a virtue, but it does mean the prior $p_{\#}(\theta)$ is not invariant to repetition of the experiment.

While it does not maximize the information $I(X; \Theta)$, the prior $p_{\#}(\theta)$ usually does very well on this measure. Figure 12 above shows an estimate of this, compared to $p_{\star}(\theta)$ as well as Jeffreys prior, and a log-normal prior.

6. Continuum limits in physics, and the renormalization group

Common to many areas of science are models in which key properties are emergent—not easily discernible from the properties of microscopic components. In many cases, a microscopically motivated model might naturally contain a large number of parameters, while the emergent behavior can be described with a much more compact effective theory. Conceptually, most parameter combinations of the underlying model are *irrelevant* for their emergent behavior. An effective theory can suffice provided it is sufficiently flexible to capture the smaller number of *relevant* parameters in the microscopic description.

These ideas are most precise for models which are *renormalizable* [96, 97]. In renormalization group (RG) analysis, a microscopic model is coarse-grained by iteratively writing down an approximate description for the behavior of a subset of the full microscopic degrees of freedom. For example, in real space RG under decimation, for a Hamiltonian which defines the probability distribution of spin states on a square lattice, a renormalized Hamiltonian should give the ‘effective Hamiltonian’ for the subspace of lattice sites for which i and j are even. Importantly, this Hamiltonian lives in the same parameter space as the microscopic one, though in general the renormalized Hamiltonian will have different parameters. This coarsening, (together with a rescaling) defines a change of parameters, and is known as an RG step; the movement of parameters under repeated iterations is known as an RG flow.

Emergent behavior is seen by considering the long-time behavior of RG flow after which models sit nearby to RG fixed points. These fixed points define possible *universality classes* of coarse-grained behavior. Phase transitions, and

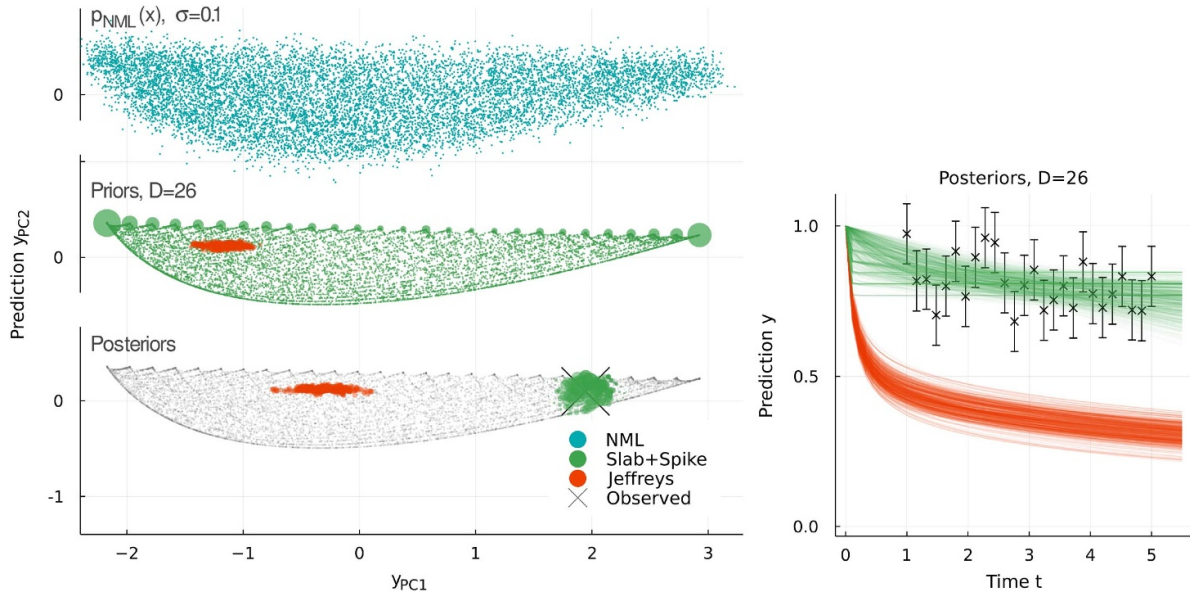


Figure 13. Behavior of the adaptive slab-and-spike prior $p_{\#}(\theta)$, for the same $D = 26$ system as in figure 11, with Jeffreys prior repeated for comparison. The distribution $p_{\text{NML}}(x)$ which when projected onto the model manifold gives the prior $p_{\#}(\theta)$ is drawn shifted upwards on the plot. As before, the posteriors in y -space are drawn shifted down on the plot. The panel on the right shows the posteriors as functions of time.

other emergent phenomena sit at nontrivial fixed points. A linearization of the flow equations close to these non-trivial fixed points has a small number of directions which are unstable, while the rest are stable. Each of these unstable directions is a relevant parameter—a small change in microscopic parameters will lead to a large change in the emergent behavior, seen after many rounds of coarse-graining. Conversely, stable directions of the flow are irrelevant—moving along them in the microscopic model will have an exponentially decreasing effect on the parameters of the emergent description. The Ising critical point has two relevant parameters (magnetic field h and reduced temperature, $t = (T - T_c)/T_c$), which are measured by the projection of the parameters onto the left eigenvectors of the linearization of RG dynamics near the Ising fixed point.

6.1. Coarse-graining and the Fisher metric

Because many models are microscopically defined, but are mostly observed through their emergent properties we sought to understand how the model manifold of microscopic parameters contracts when observables are coarse-grained. Rather than coarse-graining the parameters to directly build an effective description as is done in a usual RG treatment, we instead wanted to understand how coarse-graining changes one's ability to measure microscopic parameters [38, 39]. To do this we consider a continuous coarse-graining whereas the smallest length scale $l = l_0 \exp b$ changes the parameters change according to $d\theta^\mu/db = \beta^\mu$ where β are the beta functions that define the flow. Infinitesimal distances after such a flow can be calculated by flowing the parameters under the usual RG, and then measuring their distance in the coarse grained model, where the system size is smaller. This change in the microscopic parameter space metric can be summarized by the following equation [39]:

$$\frac{dg_{\mu\nu}}{db} = \underbrace{\beta^\alpha \partial_\alpha g_{\mu\nu}}_{g' \text{ along flow}} + \underbrace{g_{\alpha\mu} \partial_\nu \beta^\alpha + g_{\alpha\nu} \partial_\mu \beta^\alpha}_{\text{contraction of } \Delta\theta} - \underbrace{Dg_{\mu\nu}}_{\text{Vol. contraction}}. \quad (15)$$

The first two terms constitute the Lie derivative of the metric under flow β . The third term arises because the metric is extensive, proportional to the system size L^D , where D is the dimensionality of the system.

While this equation summarizes the local change in the metric, and as such the change in infinitesimal distances during coarse-graining, it remains to analyze how these distances change under typical RG flows. For this analysis it is useful to use the fact that the metric can be written as the second derivative of the free energy (in units where $\beta = 1$), $g_{\mu\nu} = \frac{\partial}{\partial \theta^\mu} \frac{\partial}{\partial \theta^\nu} F$. In renormalization group analysis it is useful to separate the free energy into a singular and an analytic piece, $F = F_a + F_s$, where the singular part is non-analytic as a function of its variables near the RG fixed point [97]. Near the RG fixed point the metric, like the free energy, can be also divided into singular and analytic pieces, $g = g^s + g^a$. By assumption, the singular contributions of the free energy are those that are preserved under RG flow—the three terms in equation (15) must precisely cancel. This means that along relevant directions—which are best measured through their influence on emergent properties, the microscopic metric does not contract¹⁴. The preservation of the metric under flow can be used to constrain the form of the metric along relevant directions [98]. However, these constraints only apply to the lower dimensional subset of parameters which are relevant in the usual RG treatment.

¹⁴ This behavior in prediction space is quite different from the flows in parameter space, where the relevant directions grow under the renormalization-group transformation.

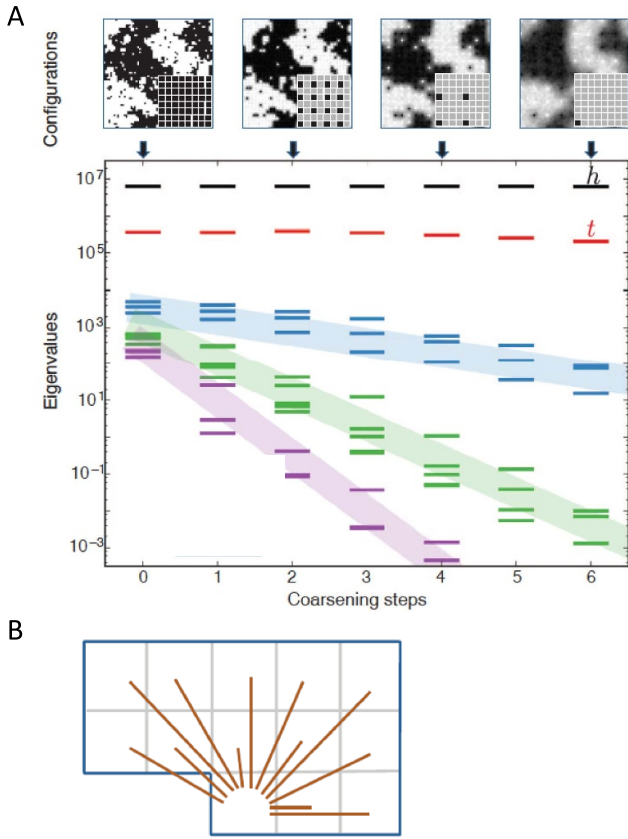


Figure 14. (A) FIM spectrum is plotted vs level of coarse-graining for a thirteen parameter ‘nearby neighbor’ model where parameters describe translation invariant couplings between spins as shown in (B) as well as a magnetic field. Without coarse-graining (0 steps), by observing entire configurations, (black in inset) relevant parameter combinations are easy to infer (red and black dashed lines) and other combinations have a typical scale. As coarse-graining proceeds (black squares in inset remain observable) relevant parameters remain easy to infer, but irrelevant ones contract with their RG exponent, slope $(-D + 2y_i)/2\log(2)$ since each coarse-graining step contracts lengths by a factor of $\sqrt{2}$. Here blue eigenvalues have $y_i = 0$, green have $y_i = -1$ and purple have $y_i = -2$. From [38]. Reprinted with permission from AAAS.

Irrelevant directions are instead measured through their influence on microscopic details and their metric is dominated by analytic contributions. For these directions the first term is expected to be small and infinitesimal distances ds contract according to:

$$\frac{d\log(ds^2)}{db} = -D + 2y_i, \quad (16)$$

where y_i is the corresponding (negative) RG exponent.

We have verified this prediction numerically for a two dimensional Ising model, where we have added couplings between non-nearest-neighbor spins (see figure 14 and [38]). We coarse-grain the system using simple decimation, only observing the red sites on a checkerboard, and so on recursively. In the Ising model these results have an intuitive

interpretation; the magnetic field and temperature can be measured easily without seeing individual spins as they control the net magnetization and the correlation length. However, other parameter combinations can only be measured by examining microscopic configurations.

6.2. Global model manifold flows

While this metric centered approach provides information about how the local structure of the manifold contracts under coarsening, it does not necessarily provide information about the global structure of the manifold. While we have fewer analytical approaches, we can numerically investigate how an entire model manifold contracts under coarse-graining. We can numerically investigate the diffusion equation, where a simple renormalization group treatment yields a single relevant parameter (the drift, v) a single marginal parameter (the diffusion constant D) and an infinite number of irrelevant parameters which control higher cumulants of the distribution [38, 39]. Numerically we investigate a hopping model, where parameters θ^μ describe the probability that a single particle moves a distance μ in a discrete time step (see figure 15(A)). We consider a seven parameter model when particles are allowed to move up to three spaces to the left and right; particle number is not necessarily conserved, but there is a six dimensional subspace which conserves particle number (the subspace where $\sum_\mu \theta^\mu = 1$).

The metric can be measured at a single point in parameter space by observing the particle density after some number of time steps. Its eigenvalue distribution is shown in figure 15(A) for the parameters where $\theta^\mu = 1/7$. In the upper panel of figure 15(A), the particle distribution is shown after t time steps. Initially the distribution is given by the parameters, $\rho_i \propto \theta^i$, and each parameter can be measured independently, yielding a metric proportional to the identity, which is not sloppy. After sufficient coarse-graining steps, the distribution approaches a Gaussian as required by the central limit theorem, characterized by a velocity and diffusion constant, and, if particle number is not conserved, by a rate of particle creation/annihilation (R). In this limit, the spectrum of the Fisher Information becomes sloppy, with the stiffest eigenvector changing R , the next stiffest changing V , then D , and with a hierarchy of irrelevant parameters which manifest only through higher cumulants in the distribution.

The non-local contraction of distances can be seen qualitatively by looking at two (two-dimensional) diffusion kernels, one a triangle, one a square, whose drift is both zero and which have the same diffusion constant (red and blue points, respectively in figure 15(B)). After a single time step these two kernels yield very clearly distinguishable particle distributions. However, after multiple time steps the two distributions become similar, approaching the same Gaussian. We numerically investigated the seven parameter model, by taking kernels uniformly from the six parameter space of particle conserving kernels, and plotting them in the first three principle components of the space of particle densities. After a single time

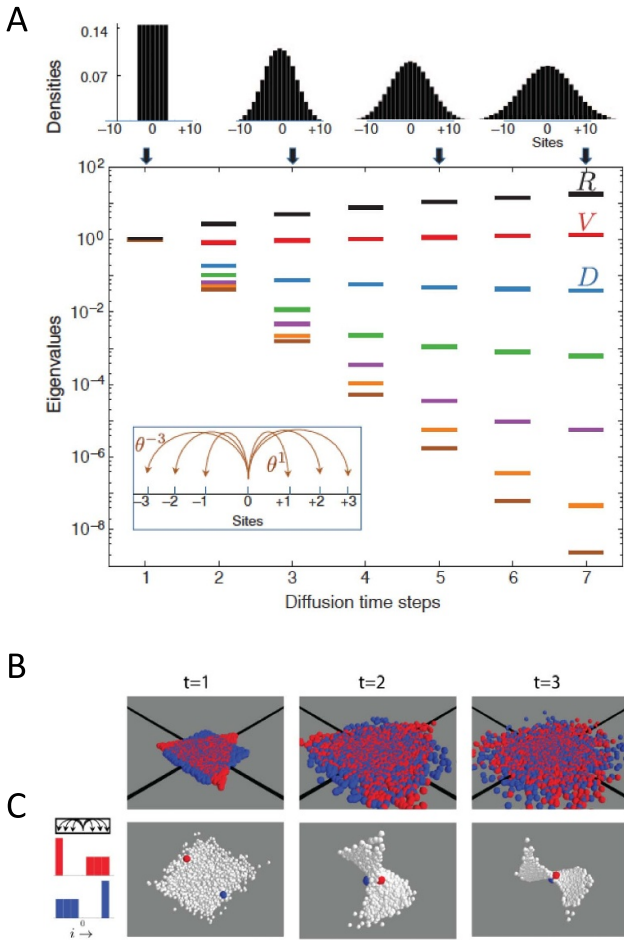


Figure 15. Coarse graining a model of hopping yields a lower dimensional diffusion equation. (A) The Fisher Information's eigenvalues for a model where particles hop according to a seven parameter kernel in each time step, with each parameter describing the rate of hopping μ sites on a discrete lattice in a discrete time step. After a single time step the metric is the identity, but if the density is observed only after taking multiple steps then the distribution becomes sloppy—the largest eigenvalue controls R , non-conservation of particle number, the next largest the drift, V , followed by the diffusion constant, D , and a hierarchy of irrelevant parameters that control higher cumulants of the density. As time evolves the central limit theorem dictates that density (above) should approach a Gaussian, parameterized by the three largest numbers, R (controlling the total number of particles), V (controlling the mean of the distribution) and D (controlling its width). Reproduced from [13]. CC BY 4.0. (B) This is shown pictorially in a model where blue particles hop according to a square kernel while red particles hop according to a triangular kernel. After a single time step their distribution is not similar. However, because their diffusion constant and drift velocity are matched, the particles' distributions become similar as time evolves. From [38]. Reprinted with permission from AAAS. (C) In the one-parameter example above, kernels (parameter sets) are drawn uniformly from the space that conserves particle number, and their model manifold's projection onto a three parameter subspace is shown at three time points. After a single time point the model manifold is not a hyper-ribbon, but is extended roughly equally in all directions. Two parameter values, marked in red and blue, are separated substantially even though they have the same values of V and D . After coarsening the model manifold collapses towards a two-dimensional surface which could be parameterized by drift and diffusion alone, and the two marked points become close to each other. Reprinted figure with permission from [39], Copyright (2018) by the American Physical Society.

step the points fall in a uniform three dimensional blob—the corresponding model manifold does not have a hyper-ribbon structure. The red and blue points, which each have zero drift and the same diffusion constant, are far apart in this model manifold. However, as time evolves the three dimensional blob collapses onto a two dimensional structure parameterized by a drift and diffusion constant. As this limit is approached the red and blue points approach each other rapidly.

In renormalizable models, after coarse-graining the model manifold contracts so that every point is close to a smaller dimensional surface which can often be described by a reduced model with just relevant parameters kept. This gives an information theoretic perspective for one of the central practical findings of the renormalization group that a simple 'toy model' will often suffice to describe a microscopically complicated system. Many systems are not renormalizable; however, they still often show a sloppy spectrum in their Fisher Information's eigenvalues and in the hyper-ribbon structure of their model manifolds. These systems would often not be sloppy if they were observed at some suitable microscopic scale. However, for macroscopic behavior, only a subset of their microscopic properties remain important, leading to local sloppiness and a global hyper-ribbon structure. In that sense the sloppiness that arises in renormalizable systems may be qualitatively similar to that arising in other systems. One key difference is that renormalizable models have a self-similarity during coarse-graining, allowing metric flow to be written compactly in terms of parameter space point flow (as in equation (15) from [39], see [99] for an interesting perspective).

7. Replicas and Minkowski embeddings for visualizing models

An essential component of working with model manifolds and hyperribbons is being able to map and visualize them. This is often achieved by sampling the manifold in a region of interest (see section 5), and plotting the resulting data set. However, finding an accurate way of visualizing complex, high-dimensional data is a difficult task. It is often not possible to faithfully represent both distance and geometry with fewer dimensions than the original data, and so it would seem that visualizing model manifolds may also not be possible. For least-squared models, we have shown both empirically (figure 4) and rigorously (section 3) that the model predictions lie on flat hyperribbons, which makes PCA visualization an effective tool for data formed by varying parameters in the model (e.g. figure 11).

Other physical models do not so neatly conform to this approach. For instance, the canonical Ising model from statistical physics predicts the likelihood of a certain spin configuration given fields and bond strengths. This does not fit the structure of a least-squares model, since the probability distributions are not Gaussian fluctuations about predicted data points, but instead a probability distribution over a data set of spin configurations. The Λ -CDM cosmological model of the early Universe similarly provides a prediction for the

likelihood of a given sky map: the probability density of finding a pattern of temperature and polarization fluctuations in the cosmic microwave background radiation. A new approach is therefore necessary to visualize these kinds of general probabilistic model manifolds, which we shall for simplicity refer to as probabilistic models in this section.

The task of visualizing high-dimensional data is well known in the data science and machine learning communities, and has fueled a growing number of visualization techniques to extract lower-dimensional visualizations that faithfully capture important and relevant properties of the manifold that are of interest to researchers. Examples include t-SNE [100], diffusion maps [101], UMAP [102], Isomap [103] and hyperbolic mapping [104]; there are many reviews [105] and software environments [106, section 2.2] to guide newcomers. In particular, all methods must address in some way the *curse of dimensionality*, where distance measures in ever increasing dimensional spaces lose useful distinguishability, mainly because the volume increases so rapidly that finite data sets become sparse and orthogonal. As a result, all visualization methods must compromise in some way. Most techniques accomplish this through a trade-off between global and local structure. For instance, classical manifold learning methods, such as PCA [107, 108] and multidimensional scaling (MDS) [109] preserve global properties at the expense of often obscuring local features. Conversely, most modern nonlinear methods, such as t-distributed stochastic network embedding (t-SNE) [100], diffusion maps [101], and hyperbolic mapping [104], highlight local structures while only qualitatively maintaining certain global features such as clusters.

There are similar problems for probabilistic models. Unlike least-squares, these models do not have a simple, finite-dimensional Euclidean space in which they can be easily embedded to reveal important features. If one wishes to preserve both local and global features, then a different compromise must be made. When visualizing collections of probability distributions, we have uncovered visualization methods which perform as well as PCA does for least-squares models, by embedding in a pseudo-Riemannian space (i.e. a space with both time-like and space-like components, for which $3 + 1$ -dimensional Minkowski space is the canonical example). Two particular choices are intensive PCA (InPCA [41]) and Intensive Symmetrized Kullback–Leibler Embedding (isKL [40]).

Using these methods, many physical models are represented as an isometric, low-dimensional hyperribbon embedding that quantitatively preserves both global and local features. The time-like coordinates with negative squared distance do make interpretation less intuitive, but often serve to illuminate important model features (figure 19). Both InPCA and isKL give hierarchical embeddings. However, isKL has the added advantage of producing a finite embedding for exponential families—although the data space for distributions can be infinite dimensional, isKL systematically extracts a space-like and a time-like component for each *parameter* in the model itself, in a hierarchical manner that reveals the underlying hyperribbon features.

7.1. Why Euclidean embeddings are doomed

To understand why it is that embedding in a pseudo-Riemannian space is more natural for manifolds of probabilistic models, consider the failures of a natural embedding method, using the Hellinger distance:

$$d_{\text{Hel}}^2[\theta, \tilde{\theta}] = 4 \sum_x \left(\sqrt{p(x|\theta)} - \sqrt{p(x|\tilde{\theta})} \right)^2. \quad (17)$$

Since probabilities are normalized, $1 = \sum_x p(x|\theta) = d_{\text{Hel}}^2[\theta, \theta]/4$, the vectors $\mathbf{n}_x = 2\sqrt{p(x|\theta)}$ are points on a hypersphere of radius two in a space whose dimension is the cardinality of the observations. The Hellinger distance is the Euclidean distance between these points. As θ is varied, the points form a manifold:

$$\mathbf{y}_x(\theta) = 2\sqrt{p(\mathbf{x}_i|\theta)}, \quad (18)$$

which can be visualized with PCA. The metric between nearby points on this manifold is given by the Fisher information¹⁵, it is *isometric*. The fact that it is isometric makes it a natural embedding, in the sense that distance between points represents the distinguishability of the different probability distributions (e.g. if two distributions have no overlap, then they are orthogonal from each other in this representation, since their dot product is zero).

The curse of dimensionality that hinders Hellinger can be seen easily in this framework (see [40]). As the dimension of the data \mathbf{x} increases, the distributions become more distinguishable from each other. They therefore become increasingly orthogonal, and all distances between non-identical distributions converge to a constant value as they become maximally apart on the hypersphere. As an illustration of this, consider the finite, two-dimensional Ising model in a field with nearest-neighbor coupling (figure 16). For small system sizes, two similar field and coupling strengths yield similar spin configurations and can be difficult to distinguish from one another. However, as the system size increases (and therefore the dimension of the observed data also increases) the increase in information allows one to more easily evaluate the underlying parameters. The probability distributions representing different parameters are more distinguishable from each other, and are therefore further apart in the hypersphere embedding. In the limit of infinite data, we can completely disentangle all possible parameter combinations: in other words, all distributions become orthogonal and equidistant from one another, forming a hypertetrahedron in probability space. The geometry of a regular hypertetrahedron with N vertices clearly cannot be captured with fewer than a $N - 1$ -dimensional embedding in a Euclidean space. But this does not remain true if we allow for both space and time-like coordinates.

¹⁵ Why do we take a square root? Nearby probability vectors $\mathbf{y}_x(\theta) = p(\mathbf{x}|\theta)$ have squared differences that do not correspond to the Fisher information distances. By taking the square root and multiplying the lengths by two, we fix this: the Hellinger embedding is *isometric*.




Figure 16. Hypersphere embeddings. (a) An original 2×2 Ising model sampled by sweeping through a range of possible field and coupling strengths and visualized using the Hellinger embedding with the first three extracted components of a classic PCA. (b) As the system size increases to 4×4 , the probability distributions are more distinguishable and thus more orthogonal, illustrating the curse of dimensionality. (c) This effect is simulated by taking four replicas of the 2×2 system. Reproduced with permission from [41].

7.2. Intensive embeddings and replica theory

If instead, we look at the information density rather than the total amount of information in a distribution, we can avoid this orthogonality catastrophe by reducing rather than increasing the dimensionality. To accomplish this, we can make use of replica theory. We can simulate the curse of dimensionality by considering replicas of the original system, i.e. multiple, independent samples of the original distributions. As an illustration, consider the two-dimensional Ising model with field and nearest neighbor interactions. A 2×2 spin system, visualized in this hypersphere embedding, has 2^4 dimensions, one for each possible spin configuration. By considering a range of possible field and coupling strengths, we sweep out a region of the model manifold. Figure 16(a) shows the first three components of a classic PCA of this embedding, with the curvature of the hypersphere very evident. As the system size increases to 4×4 spins, the dimensionality increases and the orthogonality problem becomes more evident. This effect is simulated by looking at four replicas of the 2×2 system, shown in figure 16(c).

Rather than increasing the system size with many replicas we decrease it by considering the limit of zero replicas. In this way, we obtain an *intensive* embedding. The distance between points becomes characterized not by the Hellinger distance but the Bhattacharyya¹⁶ divergence [43]:

¹⁶ The factor of eight is added to match the local distances to that of the Fisher Information metric.

$$d_{\text{Bhat}}^2[\theta, \tilde{\theta}] = -8 \log \left(\sum_x \sqrt{p(x|\theta)p(x|\tilde{\theta})} \right). \quad (19)$$

By working through the derivation of PCA using the Bhattacharyya divergence, which does not respect the triangle inequality, we obtain an embedding which is pseudo-Riemannian, called InPCA [41].

Empirical results from probabilistic models show that they share the hyperribbon structures we find in least squares models. In figure 4, for example, we see geometrically decreasing widths for the two-dimensional Ising model, variable-width Gaussians, CMB skymaps, and neural networks [40, 41]. Figure 18 shows the model manifold for a 2×2 Ising model, exhibiting projections along the first five InPCA components, clearly showing this hyperribbon decrease of widths.

7.3. Symmetrized Kullback–Leibler methods

The Hellinger distance and Bhattacharyya divergence belong to a larger class of divergences [46, 47] whose metric is given by the Fisher Information. The most commonly used distance measure for probability distributions, $p(x|\theta) \equiv p$ and $p(x|\tilde{\theta}) \equiv \tilde{p}$, is the Kullback–Leibler divergence [45] or relative entropy:

$$D_{\text{KL}}[\theta, \tilde{\theta}] = \sum_x p(x|\theta) \log \left(\frac{p(x|\theta)}{p(x|\tilde{\theta})} \right). \quad (20)$$

Because it is asymmetric, often researchers symmetrize [44, 45] the Kullback–Leibler divergence by considering the sum of the two permutations, which we write as the square of a distance d_{sKL} :

$$\begin{aligned} d_{\text{sKL}}^2 &= D_{\text{KL}}[\theta, \tilde{\theta}] + D_{\text{KL}}[\tilde{\theta}, \theta] \\ &= \sum_x \left(p(x|\theta) - p(x|\tilde{\theta}) \right) \log \left(\frac{p(x|\theta)}{p(x|\tilde{\theta})} \right). \end{aligned} \quad (21)$$

By combining the method of projecting data along orthogonal direction from MDS [41] (section 7.4) with the symmetric KL divergence, we can derive the isKL embedding [40]. An astonishing feature of this method is that it can map exponential families into a *finite dimensional* embedding space, with dimension twice the number of model parameters.

Exponential families [111] have probability distributions that can be expressed in the form familiar to physicists, as a partition function described by a Boltzmann-like distribution:

$$p(x|\theta) = h(x) \exp \left(\sum_i \eta_i(\theta) \phi_i(x) - A(\theta) \right), \quad (22)$$

where $\eta_i(\theta)$ is a function of model parameters and is known as the i th natural parameter, $\phi_i(x)$ is a function only of x and known as the i th sufficient statistic, $h(x)$ is a normalizing function, and $A(\theta)$ is the log partition function. By combining equations (21) and (22), we see that:

$$d_{\text{sKL}}^2 = \sum_i \left(\eta_i(\theta) - \eta_i(\tilde{\theta}) \right) \left(\langle \phi_i \rangle_{\theta} - \langle \phi_i \rangle_{\tilde{\theta}} \right), \quad (23)$$

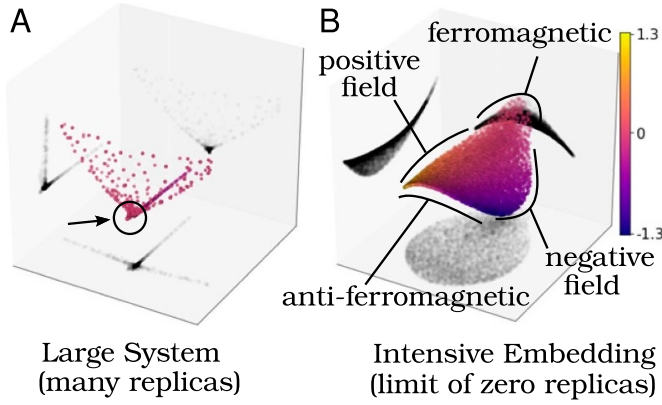


Figure 17. Intensive embedding contrasted with high-dimensional. (a) A very high-dimensional visualization of the original 2×2 Ising system by considering multiple replicas, simulating the curse of dimensionality where all points become effectively orthogonal obscuring useful features. (b) By considering the limit of zero replicas, we extract an intensive embedding (InPCA) which reveals relevant information in the model manifold (such as the ferromagnetic and anti-ferromagnetic regimes). An important difference between (a) and (b) is that the intensive embedding makes use of a pseudo-Riemannian space, with the z-component having a negative-squared distance. Reproduced with permission from [41].

where $\langle \cdot \rangle_\theta$ reflects taking an average over x using $p(x|\theta)$. Rearranging the terms in the above equation, we can see that:

$$\begin{aligned} d_{\text{sKL}}^2 &= \frac{1}{4} \sum_i \left([\eta_i(\theta) + \langle \phi_i \rangle_\theta] - [\eta_i(\tilde{\theta}) + \langle \phi_i \rangle_{\tilde{\theta}}] \right)^2 \\ &\quad - \left([\eta_i(\theta) - \langle \phi_i \rangle_\theta] - [\eta_i(\tilde{\theta}) - \langle \phi_i \rangle_{\tilde{\theta}}] \right)^2 \\ &= [\mathcal{S}_i(\theta) - \mathcal{S}_i(\tilde{\theta})]^2 - [\mathcal{T}_i(\theta) - \mathcal{T}_i(\tilde{\theta})]^2, \end{aligned} \quad (24)$$

where \mathcal{S}_i and \mathcal{T}_i represent space-like and time-like components, and can be directly calculated from the model:

$$\begin{aligned} \mathcal{S}_i(\theta) &= \frac{1}{2} (\eta_i(\theta) + \langle \phi_i \rangle_\theta), \\ \mathcal{T}_i(\theta) &= \frac{1}{2} (\eta_i(\theta) - \langle \phi_i \rangle_\theta). \end{aligned} \quad (25)$$

There are a *finite* number of components, determined by the number of parameters, regardless of the dimension of x . This provides a systematic, computationally efficient way of visualizing model manifolds, as opposed to having to calculate the distance between all points and computing the projections [40]. For example, calculating the entire probability distribution for the Ising model (needed for the Hellinger and Bhattacharyya embeddings) becomes infeasible for all but tiny system sizes (figures 16–18), but traditional Monte-Carlo methods can be used to calculate the magnetization and interaction energy needed for the isKL embedding of the Ising model for any reasonable system size (figure 19).

Moreover, isKL provides a valuable check for data without a known model. By calculating the d_{sKL}^2 between data, and visualizing both the real and imaginary components extracted using MDS, we can check to see if the resulting embedding




Figure 18. InPCA projections of the 2×2 Ising Model. The hyperribbon structure of the manifold is apparent, as the successive widths decrease geometrically (widths shown in figure 4). Color code represents field strength, matching figures 16 and 17. Orange dot represents critical point. Importantly, some components reflect a positive squared-distance (black labels) while others are a negative-squared distance apart (gray italics). Reproduced with permission from [110].

is finite or infinite. If finite, then one is led to believe that the data results from a process described by an exponential family (basically a statistical sampling described by a free energy).

7.4. Intensive embedding procedure: InPCA, isKL, and MDS

Here we discuss the methods used to generate the visualizations produced in this section. Classical PCA and MDS [41, 109] take data sets and infer features which are linearly uncorrelated. PCA allows researchers to examine data in a new basis by centering it and then extracting components of maximal variance, or diagonalize its ‘moment of inertia’, then extract projections in a hierarchical manner which minimizes the error in distances between points when restricted to a smaller dimensional subspace. Because PCA optimizes distances, one can extend it to take as inputs not the raw data but rather the distances themselves, and in this way obtain MDS. This approach can be generalized to intensive embeddings (InPCA and isKL) in the following way. First, given a set of N probability distributions \mathbf{x}_i , calculate the squared distance, i.e. the Bhattacharyya divergence or Symmetrized Kullback–Leibler divergence between them, $X_{i,j} = D(\mathbf{x}_i || \mathbf{x}_j)$ where D represents the choice of divergence. Next, generate the intensive cross-covariance matrix [41], $W = -PXP/2$, where $P_{ij} = \delta_{ij} - 1/N$ serves to mean-shift the data, so that it is centered on its center of mass.

Next, compute the eigenvalue decomposition for W , $W = U\Sigma U^T$, and note that there may be both positive and negative

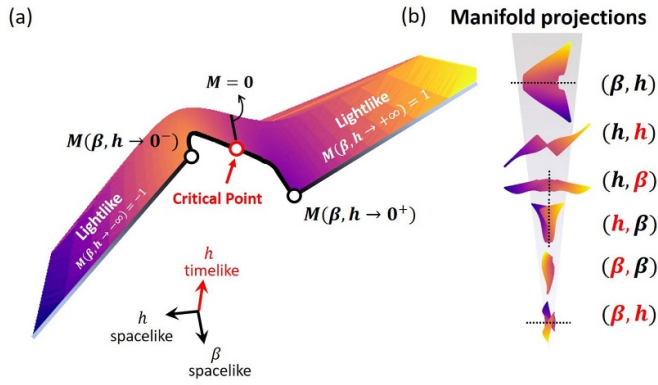


Figure 19. isKL embedding for the 2D Ising model. The coordinates $h = H/2T \pm M/2$ and $\beta = 1/2T \pm E$. H/T and $1/T$ are the natural parameters where T is the temperature and H the external field; M and E are the magnetization (total spin) and coupling energy (sum of the products of neighboring spins). Black coordinates are space-like, red are time-like (representing negative squared distance). (a) A three-dimensional image of the 4D manifold in Minkowski space. The two ‘front edges’ represent the fully up-spin and down-spin configurations, which are unintuitively spread out into a light-like line (zero net distance) as the external field varies. On the other hand, it adds intuition to the jump in the magnetization as the field crosses zero; the $\pm M$ branches are nicely separated by a light-like displacement, faithfully representing the scaling behavior at the critical point. (b) 2D manifold projections along the six pairs of coordinates. Reprinted figure with permission from [40], Copyright (2020) by the American Physical Society.

eigenvalues. It is important that they be ranked by magnitude, *i.e.* $|\Sigma_{00}| \geq |\Sigma_{11}| \geq \dots$. The coordinate projections for the intensive embedding are determined by:

$$T = U\sqrt{\Sigma}. \quad (26)$$

One then plots the projections using the columns of T . Coordinates corresponding to positive eigenvalues represent space-like components, in that they have a positive-squared contribution to the distance between points. Conversely, coordinates corresponding to *negative* eigenvalues represent time-like components, in that they have negative-squared contributions to the distance between points. The resulting figures are interpreted much like a spacetime diagram, where points with zero distance are light-like separated. Note that, because all points are a positive-squared distance apart (since their divergence is non-negative) their overall distance will be a space-like separation, although they may be visualized using both space-like and time-like components.

8. Future directions

In this review, we discussed how information geometry explains sloppiness in multiparameter models: predictions are independent of all but a few parameter combinations, which is unfortunate if you want to deduce the parameters from collective behavior, but fortunate if you want to test the predictions of the model before the parameters are fully measured. Extensive observations [68] and rigorous results [30] explain this by noting that the predictions form a hyperribbon structure. We

have described the MBAM method [14] for deriving simpler, emergent models compatible with the data using the boundaries of the model manifold along thin axes. We have introduced new Bayesian priors [36], that implement Occam’s razor by placing δ -function preferences for these simpler boundary models. We have connected the hyperribbon structure of the model manifold in emergent statistical mechanics models to renormalization-group flows [39]. And we have introduced visualization methods that bypass the curse of dimensionality for high-dimensional data sets by generating isometric hyperribbon embeddings in Minkowski-like spaces [40, 41].

What is there left to do? In previous reviews and papers, we have explored and speculated about applications of information geometry to systems biology [68], power systems [32], robustness and neutral spaces in biology [7], controlling and optimizing complex instruments such as particle accelerators [112], and explaining why science works [13]. Here we focus on possible future developments related to the new tools and methods discussed here—understanding why hyperribbons and emergence arise, using their boundaries as simpler models, relations to emergence in physics, new Bayesian priors, and visualization methods that bypass the curse of dimensionality.

8.1. Visualization, experimental design, and testing model validity

The sloppiness of multiparameter models suggests that extracting parameters from data will be challenging or impossible. Model reduction methods such as MBAM show that one can describe data in complex systems with emergent models for which many parameters and model components are simply ignored (set to zero, infinity, or replaced by other parameters), *even when those components of the model are microscopically known to be important and relevant*. While we may find the emergent models more comprehensible, are not they wrong? The rates in a system biology model do not go to infinity, the enzymes are not in the limit of short supply—the simpler model works by shifting other parameters to compensate for these simplifying assumptions.

The progress of science often exhibits this in reverse—one develops a simple theory which captures the essential features of a problem, and then adds complexity to address new experimental facts. Can we use information geometry to guide the design of new experiments to pin down models in an optimal way? Optimal experimental design [113] is a well-developed branch of statistics, but usually is done in a linearized approximation (but see [114])—can we use hyperribbon embeddings and differential geometry to turn experimental design into a more global tool? As a particular example, the InPCA methods were developed in order to visualize the space of possible Universes as reflected by current measurements of the microwave background radiation. The next generation of instruments is designed to measure and test new physical phenomena—gravitational waves from the epoch of inflation, evidence of light relic particles in the early Universe, and signatures of dark matter in the cosmic microwave background radiation. Would it be useful to visualize the predictions of

extensions of the Λ CDM model that include these phenomena, by mapping out those portions of the model manifold that are consistent with current experimental data? Can we use such visualizations to help guide the choice of the next generation of instruments?

Our information geometry gives us a theoretical justification for trusting the predictions of nonlinear models with many undetermined parameters. Climate science, for example, incorporates not only well-controlled theories of heat, hydrodynamics, and radiation, but much more speculative models of clouds, glaciers, and forest ecologies (not to mention human behavior). It is reassuring to hear that equally complex models in systems biology with even more undetermined parameters are quite useful in guiding drug design (see below). But just because complex models fitted to data can make correct predictions does not mean they must do so. Statistics provides many tests of model validity and many methods to compare the predictive power of rival models, for cases where the parameters are well determined. Can we use information geometry methods to create similar tests and comparison tools for multiparameter models where the uncertainties in model parameters are typically not only outside the linear regime, but often unbounded?

8.2. Understanding emergent low-dimensional behaviors

Emergence in physics relies on small parameters. Continuum limits work (elasticity, fluids, superfluids, etc) when frequencies and wavelengths are low; the renormalization group works (emergent scale invariance, avalanches, crackling noise, etc) when in addition there is a proximity to a continuous transition between states. Sections 6 and 7.3 discussed how information geometry applied to these models form hyperribbons. On the other hand, we used interpolation theory to explain the low-dimensional emergent behavior and reduced models that arise in nonlinear least-squares models fit to data (section 3). Emergence in physics and emergence in least-squares models are both characterized by sloppiness (only a few combinations of microscopic parameters dominate the behavior), and both form hyperribbons, but our explanations for the two seem completely different. For more general probabilistic models, intensive embeddings (section 7) usually lead to hyperribbons and emergent theories, but we do not yet have an explanation for this in general, except for the special (but very useful) case of exponential families, where the model manifold not only is thinner and thinner along successive cross sections, but becomes zero thickness after a finite number of embedding dimensions.

Can we find a common theme between these disparate results—a unified theory for the emergence of simplicity from complexity? One possible clue is that sloppiness is a property both of the scientific model and the predictions of interest. The Ising model is not sloppy, and has no beautiful emergent theory, unless one only cares about long length and time scales. And a systems biology model is not sloppy if one designs experiments to measure one parameter at a time. Multiparameter least squares models have emergent simplicity when one focuses on collective behavior of the system as a whole,

and statistical mechanics has emergent simplicity when one focuses on collective behavior of large numbers of particles. Perhaps it is caring about the answers to the right questions that makes science possible?

8.3. Intensive embeddings for big data, from biology to deep neural networks

In our work visualizing the model manifold for probabilistic models (section 7), we grappled with the *curse of dimensionality*, commonplace in big data applications. We identified the problem of having too much information: data points in too high a dimensional space naturally becoming orthogonal. We cured the problem first using replica theory and then using an analytic method for models that form ‘exponential families’. Our cures work because of two features. (a) We found a way of embedding our model in a Minkowski-like space, using MDS ([41], section 7.4), that violates the triangle inequality. (b) We use measures of squared distance between data points (Bhattacharyya and symmetrized Kullback–Leibler) that diverge logarithmically with the overlap between the probability distributions. (This ensures that doubling the amount of data simply rescales all distances by a factor of 2).

Can we apply these basic ideas to other big data problems, that are not framed as probabilistic models? We know of at least one machine learning method, PHATE [115] (due to Kevin Moon) that has applied these two principles to visualize biological data (in particular single-cell RNA sequencing data sets). Here let us consider two possible applications to machine learning, as represented by the millions-of-parameter deep learning neural network models.

Deep neural networks are ‘big’ in the number of neurons per layer (hundreds to thousands), the number of neural connections between layers (millions of weight parameters), the number of dimensions in the initial data used to train the network (say, one per pixel in the images being recognized), and the number of data points used to train the network (millions of images). They have been remarkably successful at extracting meaningful information from these large data sets, suggesting that there is some hidden, low-dimensional structures [116] in the data, that is disentangled and flattened via the nonlinear transformations through succeeding neural layers.

With Pratik Chaudhari, some of us have been exploring [117] how one might use intensive embeddings to visualize both the training of these neural networks (as in [41]) and potentially explore low-dimensional structures discovered after the training process. Consider a network that categorizes, taking N images \mathbf{x} and returns a preference probability $\rho(\mathbf{x})$ whose largest value is the predicted category. The vector $\mathbf{R} = (1/N)\{\rho(x_1), \rho(x_2), \dots\}$ can be viewed as a probability distribution over the whole set, and one could use InPCA or isKL to visualize how the network trains with time, averaging over the random initial states and training steps. Visualization of the resulting structures for training and test data could lend geometrical insight into how and why different training protocols differ in their generalization errors. Alternatively, one could examine the neural activations $h(\mathbf{x})$ of the different layers of a single neural network—each encoding distilled

information about the initial data \mathbf{x} gleaned from the previous layer that is most relevant to the task at hand. Using PCA on the vectors $h(\mathbf{x})$ is likely doomed for the standard reasons, but conceivably using MDS with an intensive metric could allow one to see the emergence of a low dimensional (perhaps hyperribbon-like) representation of the data as one passes through successive trained layers.

8.4. Possible applications of model reduction

The manifold boundary approximation method is a principled way to select effective theories for sloppy systems. MBAM is objective, data-driven, and does not require expert intuition. Through sequences of simplifying approximations, it gives identifiable models that retain essential elements, such as dominant feedback loops, while remaining expressed in terms of microscopic physics. We anticipate that simplified models will be increasingly important by enabling clearer reasoning about big data and complex processes.

Recent result from systems biology modeling the Wnt pathway illustrate how this may be done. Wnt signaling is the primary mechanism by which short-range extra-cellular signals are relayed to the nucleus, and it plays a major role in normal embryonic development as well as cancer tumorigenesis. While the real system is notoriously complex, it regularly exhibits simple, comprehensible behavior. For example, the steady state conditions of the pathway may encode for the level of extra-cellular Wnt [118] while at other times, a negative feedback drives sustained oscillations that regulate somitogenesis [119]. In each case, the behavior is explicable in terms of simple models with only a few basic mechanisms. However, by reasoning about the relationships between these reduced models within the larger space of candidate approximations, more is possible. The *supremal model*, i.e. simplest model reducible to each effective behavior includes the relevant mechanisms for *transitioning* between states [120]. This model makes predictions that are qualitatively different from the data to which it was trained and is a starting point for controller design.

Can simple, effective models facilitate the transition of big data into practical knowledge in spite of being incomplete? Quantitative systems pharmacology uses complex models to design and engineer drugs [121], but reasoning about causality within complex models is difficult. In our conversations with the FDA and pharmaceutical companies, challenges in justifying new drugs and human trials are closely related to model complexity [122]. If it is unreasonable to identify all of the parameters in realistic biological models, how can drug developers meet standards of government approval agencies [123]? Simple models that are rigorously derived from first principles can be more easily and cheaply validated. Furthermore, simple models that contain only the relevant mechanism are more readily communicated and trusted in ways that complicated simulations are not.

Can similar reasoning be applied beyond biology? Empirical potentials in molecular modeling are designed to mimic

quantum effects within a classical framework. But which combination of quantum effects are responsible for specific macroscopic material properties? The combinatorial large space of stoichiometry and material properties make the problem particularly challenging. Large databases of quantum calculations [124, 125] archive data while projects such as OpenKIM [126] satisfy the dual problem of documenting and archiving useful models. But leveraging these data and models to design new materials remains an open problem. Can simple models that are successful in a limited context help clarify what must be accurately modeled to guide material by design?

We have speculated about the relationship between sloppiness and large machine learning models. Although model size seems to play an important role in learning capacity and training rate, large models cannot be deployed on small devices without a reduction step. Furthermore, large black-box models also do not lead to explicit, declarative knowledge in the same way as minimal, parsimonious representation [127]. The nascent fields of explainable AI [128] and physics-guided machine learning [129–131] may also benefit from simplification in a context-adaptive way.

We have also mentioned climate as a field whose predictions are the result of multi-physics simulations. Uncertainties in specific components (e.g. models of clouds or forest ecology), as well as their role in the larger simulation, make it difficult to reason about causality in such simulations. Simplified models that automatically identify the essential feedback loops and effective control knobs may guide how modelers understand how these uncertain components affect conclusions. They may also aid science communication and education for policy makers and the general public.

The situation in climate is reminiscent of that in power systems that similarly use multi-physics to model large engineered systems. When complex models fail to match data (even presumably identified models), for example in well-instrumented blackouts [132], it is bewildering to identify the source of the problem. Simplified models make it easier to identify and correct errors in the model. Our model reduction methods have made inroads towards this end in power systems [32]; can they similarly find use in these other areas?

Power systems application also face the challenge that they operate at the intersection of industry and government where issues of individual privacy, proprietary information, and public security are relevant. Model reduction is a means to limit and regulate how information is shared or redacted among vested parties. Models that contain sensitive infrastructure information can be systematically and objectively coarse-grained to remove classified pieces while simultaneously meeting the demands of public disclosure and oversight.

The computational cost of solving the geodesic equation limit the scope of current algorithms to models with at most a few hundred parameters, so alternative strategies are needed. It may be possible to guide optimization algorithms toward the boundary; merging the fitting and reduction steps. Alternative approaches may focus on targeted subsets of parameters. The topological relationship among potential models remains

largely unexplored, but such a theory could better inform model selection for target applications.

8.5. Conclusions

Efficiently reasoning with big models and big data will require new organizational paradigms [133, 134]. The physical theories of the twentieth century, interconnected through a web of classical limits as in figure 8, suggest a way forward. Effective theories organize our understanding and focus our attention to the relevant physical principles within specific regimes of validity. The hyperribbons presented here reflect a similar information hierarchy, but sloppiness adds a new wrinkle to the traditional narrative by justifying simple theories without resorting to classical limits, coarse-graining, or renormalization. Visualization tools and model reduction methods extract simple explanations in a context-specific way, and our theories of information topology organize them into a network of interrelated models, each applicable at an appropriate scale and for specific applications. While science has traditionally progressed by discovering effective theories before the underlying mechanism, the future of science may well operate in the opposite direction, cataloging and organizing effective theories for the Universe's myriad complexities. Sloppy models abetted with the formalism of information geometry explain why and how this can be done.

Data availability statement

No new data were created or analyzed in this study.

Acknowledgments

We would like to thank David Schwab and Pratik Chaudhari for helpful conversations. K N Q was supported in part by the US National Science Foundation, through the Center for the Physics of Biological Function (PHY-1734030) and from NSF DMR-1719490. Parts of this work were performed at the Aspen Center for Physics, supported by NSF Grant No. PHY-1607611. The participation of M C A at Aspen was supported by the Simons Foundation. M K T was supported by NSF EPCN-1710727, DMR-1753357, DMR-1834332, and DMR-1834251, NSF CMMT-1753357, and NSF CMMT-1834332. J P S was supported by NSF DMR-1719490. B B M and M C A were supported by a Simons Investigator award, and B B M was supported by NIH R35 GM138341.

ORCID iDs

Katherine N Quinn  <https://orcid.org/0000-0002-8601-2912>

Michael C Abbott  <https://orcid.org/0000-0003-1091-6357>

Mark K Transtrum  <https://orcid.org/0000-0001-9529-9399>

Benjamin B Machta  <https://orcid.org/0000-0001-9065-614X>

James P Sethna  <https://orcid.org/0000-0001-9126-0892>

References

- [1] Brown K S and Sethna J P 2003 Statistical mechanical approaches to models with many poorly known parameters *Phys. Rev. E* **68** 021904
- [2] Brown K S, Hill C C, Calero G A, Myers C R, Lee K H, Sethna J P and Cerione R A 2004 The statistical mechanics of complex signaling networks: nerve growth factor signaling *Phys. Biol.* **1** 184–95
- [3] Amari S-I and Nagaoka H 2000 *Methods of Information Geometry* vol 191 (Providence, RI: American Mathematical Society)
- [4] Amari S-I 2016 Information geometry and its applications *Applied Mathematical Sciences* vol 194 (Tokyo: Springer)
- [5] Nielsen F 2020–10 An elementary introduction to information geometry *Entropy* **22** 1100
- [6] Nielsen F 2022 The many faces of information geometry *Not. Am. Math. Soc.* **69** 36–45
- [7] Daniels B C, Chen Y-J, Sethna J P, Gutenkunst R N and Myers C R 2008 Sloppiness, robustness and evolvability in systems biology *Curr. Opin. Biotechnol.* **19** 389–95
- [8] Draghi J A, Parsons T L, Wagner G P and Plotkin J B 2010 Mutational robustness can facilitate adaptation *Nature* **463** 353–5
- [9] Tian T, Olson S, Whitacre J M and Harding A 2011 The origins of cancer robustness and evolvability *Integr. Biol.* **3** 17–30
- [10] O'Leary T, Williams A H, Caplan J S and Marder E 2013 Correlations in ion channel expression emerge from homeostatic tuning rules *Proc. Natl Acad. Sci.* **110** E2645–54
- [11] O'Leary T, Sutton A C and Marder E 2015 Computational models in the age of large datasets *Curr. Opin. Neurobiol.* **32** 87–94
- [12] Ori H, Marder E and Marom S 2018 Cellular function given parametric variation in the Hodgkin and Huxley model of excitability *Proc. Natl Acad. Sci.* **115** E8211–8
- [13] 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 *J. Chem. Phys.* **143** 010901
- [14] Transtrum M K and Qiu P 2014 Model reduction by manifold boundaries *Phys. Rev. Lett.* **113** 098701
- [15] Constantine P G 2015 *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies* (Philadelphia, PA: SIAM)
- [16] Maiwald T *et al* 2016 Driving the model to its limit: profile likelihood based model reduction *PLoS One* **11** e0162366
- [17] Fröhlich F *et al* 2018 Efficient parameter estimation enables the prediction of drug response using a mechanistic pan-cancer pathway model *Cell Syst.* **7** 567–79.e6
- [18] Holiday A, Kooshkbaghi M, Bello-Rivas J M, William Gear C, Zagaris A and Kevrekidis I G 2019 Manifold learning for parameter reduction *J. Comput. Phys.* **392** 419–31
- [19] Frankle J and Carbin M 2019 The lottery ticket hypothesis: finding sparse, trainable neural networks (arXiv:1803.03635)
- [20] Raman D V, Anderson J and Papachristodoulou A 2017 Delineating parameter unidentifiabilities in complex models *Phys. Rev. E* **95** 032314
- [21] Brouwer A F and Eisenberg M C 2018 The underlying connections between identifiability, active subspaces, and parameter space dimension reduction (arXiv:1802.05641)
- [22] Chis O-T, Villaverde A F, Banga J R and Balsa-Canto E 2016 On the relationship between sloppiness and identifiability *Math. Biosci.* **282** 147–61
- [23] Dufresne E, Harrington H A and Raman D V 2018 The geometry of sloppiness *J. Algebr. Stat.* **9** 30–68

- [24] Apgar J F, Witmer D K, White F M and Tidor B 2010 Sloppy models, parameter uncertainty and the role of experimental design *Mol. BioSyst.* **6** 1890
- [25] Lill D, Timmer J and Kaschek D 2019 Local Riemannian geometry of model manifolds and its implications for practical parameter identifiability *PLoS One* **14** e0217837
- [26] LaMont C H and Wiggins P A 2019 Correspondence between thermodynamics and inference *Phys. Rev. E* **99** 052140
- [27] Schwartz G 1978 Estimating the dimension of a model *Ann. Stat.* **6** 461–4
- [28] Akaike H 1998 Information theory and an extension of the maximum likelihood principle *Selected Papers of Hirotugu Akaike* ed E Parzen, K Tanabe and G Kitagawa (New York: Springer) pp 199–213
- [29] Kaplan J, McCandlish S, Henighan T, Brown T B, Chess B, Child R, Gray S, Radford A, Wu J and Amodei D 2020 Scaling laws for neural language models (arXiv:2001.08361)
- [30] Quinn K N, Wilber H, Townsend A and Sethna J P 2019 Chebyshev approximation and the global geometry of model predictions *Phys. Rev. Lett.* **122** 158302
- [31] Transtrum M K and Qiu P 2016 Bridging mechanistic and phenomenological models of complex biological systems *PLoS Comput. Biol.* **12** e1004915–34
- [32] Transtrum M K, Saric A T and Stankovic A M 2017 Measurement-directed reduction of dynamic models in power systems *IEEE Trans. Power Syst.* **32** 2243–53
- [33] Niksic T and Vretenar D 2016 Sloppy nuclear energy density functionals: effective model reduction *Phys. Rev. C* **94** 024333
- [34] Böhner G and Venkataraman G 2017 Identifiability, reducibility and adaptability in allosteric macromolecules *J. Gen. Physiol.* **149** 547–60
- [35] Lombardo D M and Rappel W-J 2017 Systematic reduction of a detailed atrial myocyte model *Chaos* **27** 093914
- [36] Mattingly H H, Transtrum M K, Abbott M C and Machta B B 2018 Maximizing the information learned from finite data selects a simple model *Proc. Natl Acad. Sci.* **115** 1760–5
- [37] Abbott M C and Machta B B 2022 Far from asymptopia: unbiased high-dimensional inference cannot assume unlimited data (arXiv:2205.03343)
- [38] Machta B B, Chachra R, Transtrum M K and Sethna J P 2013 Parameter space compression underlies emergent theories and predictive models *Science* **342** 604–7
- [39] Raju A, Machta B B and Sethna J P 2018 Information loss under coarse graining: a geometric approach *Phys. Rev. E* **98** 052112
- [40] Teoh H K, Quinn K N, Kent-Dobias J, Clement C B, Xu Q and Sethna J P 2020 Visualizing probabilistic models in Minkowski space with intensive symmetrized Kullback–Leibler embedding *Phys. Rev. Res.* **2** 1–16
- [41] Quinn K N, Clement C B, De Bernardis F, Niemack M D and Sethna J P 2019 Visualizing probabilistic models and data with intensive principal component analysis *Proc. Natl Acad. Sci.* **116** 13762–7
- [42] Hellinger E 1909 Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen *J. Reine Angew. Math.* **1909** 210–71
- [43] Bhattacharyya A 1946 On a measure of divergence between two multinomial populations *Sankhyā* **7** 401–6
- [44] Jeffreys H 1939 *Theory of Probability* 2nd edn (Oxford: Clarendon)
- [45] Kullback S and Leibler R A 1951 On information and sufficiency *Ann. Math. Stat.* **22** 79–86
- [46] Csizsár I and Shields P C *et al* 2004 Information theory and statistics: a tutorial *Found. Trends Commun. Inf. Theory* **1** 417–528
- [47] Rényi A *et al* 1961 On measures of entropy and information *Proc. 4th Berkeley Symp. on Mathematical Statistics and Probability* (Berkeley, CA) vol 1 pp 547–61
- [48] Fisher R A 1922 On the mathematical foundations of theoretical statistics *Phil. Trans. R. Soc. A* **222** 309–68
- [49] Rao C R 1945 Information and the accuracy attainable in the estimation of statistical parameters *Reson. J. Sci. Educ.* **20** 78–90
- [50] Cramér H 1946 *Mathematical Methods of Statistics* (PMS-9) vol 9 (Princeton, NJ: Princeton University Press)
- [51] Rao C R 1992 *Information and the Accuracy Attainable in the Estimation of Statistical Parameters* (New York: Springer)
- [52] Sethna J P *et al* 2017 Deformation of crystals: connections with statistical physics *Annu. Rev. Mater. Res.* **47** 217–46
- [53] Transtrum M K, Machta B B and Sethna J P 2011 Geometry of nonlinear least squares with applications to sloppy models and optimization *Phys. Rev. E* **83** 036701
- [54] Waterfall J J, Casey F P, Gutenkunst R N, Brown K S, Myers C R, Brouwer P W, Elser V and Sethna J P 2006 Sloppy-model universality class and the vandermonde matrix *Phys. Rev. Lett.* **97** 150601
- [55] Tofteskov J, Tørngren M A, Bailey N P and Hansen J S 2019 Modelling headspace dynamics in modified atmosphere packaged meat *J. Food Eng.* **248** 46–52
- [56] Gutenkunst R 2007 Sloppiness, modeling, and evolution in biochemical networks *PhD Dissertation* Cornell University
- [57] Chachra R, Transtrum M K and Sethna J P 2012 Structural susceptibility and separation of time scales in the van der Pol oscillator *Phys. Rev. E* **86** 026712
- [58] Wen M, Shirodkar S N, Plecháč P, Kaxiras E, Elliott R S and Tadmor E B 2017 A force-matching Stillinger–Weber potential for MoS₂: parameterization and Fisher information theory based sensitivity analysis *J. Appl. Phys.* **122** 244301
- [59] Transtrum M K, Francis B L, Saric A T and Stankovic A M 2018 Simultaneous global identification of dynamic and network parameters in transient stability studies 2018 *IEEE Power & Energy Society General Meeting (PESGM)* (IEEE) pp 1–5
- [60] Daniels B C and Nemenman I 2015 Automated adaptive inference of phenomenological dynamical models *Nat. Commun.* **6** 8133
- [61] Mortenson M, Nielsen T, Transtrum M K and Knoble D 2022 Parameter inference in seabed acoustic models JASA paper *J. Theor. Comput. Acoust.* (accepted)
- [62] Pepiot P 2018 private communication
- [63] Transtrum M K, Machta B B and Sethna J P 2010 Why are nonlinear fits to data so challenging? *Phys. Rev. Lett.* **104** 060201
- [64] Transtrum M K, Machta B B and Sethna J P 2011 Geometry of nonlinear least squares with applications to sloppy models and optimization *Phys. Rev. E* **83** 036701
- [65] Apgar J F, Witmer D K, White F M and Tidor B 2010 Sloppy models, parameter uncertainty and the role of experimental design *Mol. BioSyst.* **6** 1890–900
- [66] Chachra R, Transtrum M K and Sethna J P 2011 Comment on “Sloppy models, parameter uncertainty and the role of experimental design” *Mol. BioSyst.* **7** 2522
- [67] Zwolak J W, Tyson J J and Watson L T 2005 Globally optimised parameters for a model of mitotic control in frog egg extracts *IEE Proc.-Syst. Biol.* **152** 81–92
- [68] 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 Comput. Biol.* **3** e189–8

- [69] Holiday A, Kooshkbaghi M, Bello-Rivas J M, William Gear C, Zagaris A and Kevrekidis I G 2019 Manifold learning for parameter reduction *J. Comput. Phys.* **392** 419–31
- [70] Gorelik G E 2005 Matvei bronstein and quantum gravity: 70th anniversary of the unsolved problem *Phys.-Usp.* **48** 1039
- [71] Kuchař K 1980 Gravitation, geometry and nonrelativistic quantum theory *Phys. Rev. D* **22** 1285
- [72] Padmanabhan H and Padmanabhan T 2011 Nonrelativistic limit of quantum field theory in inertial and noninertial frames and the principle of equivalence *Phys. Rev. D* **84** 085018
- [73] Tegmark M 2008 The mathematical universe *Found. Phys.* **38** 101–50
- [74] Eliashberg G 1960 Interactions between electrons and lattice vibrations in a superconductor *Sov. Phys.-JETP* **11** 696–702
- [75] Gor'kov L 1958 On the energy spectrum of superconductors *Sov. Phys.-JETP* **7** 158
- [76] Eilenberger G 1968 Transformation of Gorkov's equation for type II superconductors into transport-like equations *Z. Phys. A* **214** 195–213
- [77] Gor'kov L P 1959 Microscopic derivation of the Ginzburg–Landau equations in the theory of superconductivity *Sov. Phys.-JETP* **9** 1364–7
- [78] Transtrum M K 2016 Manifold boundaries give gray-box approximations of complex models (arXiv:1605.08705)
- [79] Transtrum M K, Hart G and Qiu P 2014 Information topology identifies emergent model classes (arXiv:1409.6203)
- [80] Petzold L 1982 Differential-algebraic equations are not ode's *SIAM J. Sci. Stat. Comput.* **3** 367–84
- [81] Kass R E and Wasserman L 1996 The selection of prior distributions by formal rules *J. Am. Stat. Assoc.* **91** 1343–70
- [82] Jeffreys H 1946 An invariant form for the prior probability in estimation problems *Proc. R. Soc. A* **186** 453–61
- [83] Lindley D V 1956 On a measure of the information provided by an experiment *Ann. Math. Stat.* **27** 986–1005
- [84] Bernardo J M 1979 Reference posterior distributions for Bayesian-inference *J. R. Stat. Soc. B* **41** 113–47
- [85] Clarke B S and Barron A R 1994 Jeffreys' prior is asymptotically least favorable under entropy risk *J. Stat. Plan. Inference* **41** 37–60
- [86] Färber G 1967 Die Kanalkapazität allgemeiner Übertragungskkanäle bei begrenztem Signalwertbereich beliebigen Signalübertragungszeiten sowie beliebiger Störung *Arch. Elektr. Uebertrag.* **21** 565–74
- [87] Smith J G 1971 The information capacity of amplitude- and variance-constrained scalar Gaussian channels *Inf. Control* **18** 203–19
- [88] Berger J O, Bernardo J M and Mendoza M 1989 On priors that maximize expected information *Recent Developments in Statistics and Their Applications* (Seoul: Freedom Academy Publishing)
- [89] Sims C A 2003 Implications of rational inattention *J. Monet. Econ.* **50** 665–90
- [90] Abbott M C and Machta B B 2019 A scaling law from discrete to continuous solutions of channel capacity problems in the low-noise limit *J. Stat. Phys.* **176** 214–27
- [91] Aoyama T, Kinoshita T and Nio M 2019 Theory of the anomalous magnetic moment of the electron *Atoms* **7** 28
- [92] Huber M F, Bailey T, Durrant-Whyte H and Hanebeck U D 2008 On entropy approximation for gaussian mixture random vectors 2008 *IEEE Int. Conf. Multisensor Fusion and Integration for Intelligent Systems* (IEEE) pp 181–8
- [93] Myung J I, Navarro D J and Pitt M A 2006 Model selection by normalized maximum likelihood *J. Math. Psychol.* **50** 167–79
- [94] Grünwald P and Roos T 2019 Minimum description length revisited *Int. J. Math. Ind.* **11** 1930001
- [95] Mitchell T J and Beauchamp J J 1988 Bayesian variable selection in linear regression *J. Am. Stat. Assoc.* **83** 1023–32
- [96] Sethna J P 2006 *Statistical Mechanics: Entropy, Order Parameters and Complexity* (Oxford: Oxford University Press)
- [97] Goldenfeld N 1972 *Lectures On Phase Transitions and the Renormalization Group* (Reading, MA: Addison-Wesley)
- [98] Maity R, Mahapatra S and Sarkar T 2015 Information geometry and the renormalization group *Phys. Rev. E* **92** 052101
- [99] Strandkvist C, Chvykov P and Tikhonov M 2020 Beyond RG: from parameter flow to metric flow (arXiv:2011.12420)
- [100] van der Maaten L and Hinton G 2008 Visualizing data using t-SNE *J. Mach. Learn. Res.* **9** 2579–605
- [101] Coifman R R, Lafon S, Lee A B, Maggioni M, Nadler B, Warner F and Zucker S W 2005 Geometric diffusions as a tool for harmonic analysis and structure definition of data: diffusion maps *Proc. Natl Acad. Sci.* **102** 7426–31
- [102] McInnes L, Healy J and Melville J 2018 UMAP: uniform manifold approximation and projection for dimension reduction (arXiv:1802.03426)
- [103] Tenenbaum J B, Silva V D and Langford J C 2000 A global geometric framework for nonlinear dimensionality reduction *Science* **290** 2319–23
- [104] Boguná M, Papadopoulos F and Krioukov D 2010 Sustaining the internet with hyperbolic mapping *Nat. Commun.* **1** 1–8
- [105] Izenman A J 2012 *Introduction to manifold learning* Wiley *Interdiscip. Rev.-Comput. Stat.* **4** 439–46
- [106] Pedregosa F et al 2011 Scikit-learn: machine learning in Python *J. Mach. Learn. Res.* **12** 2825–30
- [107] Pearson K 1901 LIII on lines and planes of closest fit to systems of points in space *London, Edinburgh Dublin Phil. Mag. J. Sci.* **2** 559–72
- [108] Hotelling H 1933 Analysis of a complex of statistical variables into principal components *J. Educ. Psychol.* **24** 417
- [109] Torgerson W S 1952 Multidimensional scaling: I. Theory and method *Psychometrika* **17** 401–19
- [110] Quinn K N 2019 Patterns of structural hierarchies in complex systems *PhD Dissertation* Cornell University
- [111] Nielsen F and Garcia V 2009 Statistical exponential families: a digest with flash cards (arXiv:0911.4863)
- [112] Bergan W F, Bazarov I V, Duncan C J R, Liarte D B, Rubin D L and Sethna J P 2019 Online storage ring optimization using dimension-reduction and genetic algorithms *Phys. Rev. Accel. Beams* **22** 054601
- [113] Casey F P, Baird D, Feng Q, Gutenkunst R N, Waterfall J J, Myers C R, Brown K S, Cerione R A and Sethna J P 2007 Optimal experimental design in an epidermal growth factor receptor signalling and down-regulation model *IET Syst. Biol.* **1** 190–202
- [114] Vanlier J, Tiemann C A, Hilbers P A and van Riel N A 2012 A Bayesian approach to targeted experiment design *Bioinformatics* **28** 1136–42
- [115] Moon K R et al 2019 Visualizing structure and transitions in high-dimensional biological data *Nat. Biotechnol.* **37** 1482–92
- [116] Goldt S, Mézard M, Krzakala F and Zdeborová L 2020 Modeling the influence of data structure on learning in neural networks: the hidden manifold model *Phys. Rev. X* **10** 041044
- [117] Ramesh R, Mao J, Griniasty I, Yang R, Teoh H K, Transtrum M, Sethna J P and Chaudhari P 2022 A picture of the space of typical learnable tasks (<https://doi.org/10.48550/arXiv.2210.17011>)

- [118] Goentoro L and Kirschner M W 2009 Evidence that fold-change and not absolute level, of β -catenin dictates wnt signaling *Mol. Cell* **36** 872–84
- [119] Jensen P B, Pedersen L, Krishna S and Jensen M H 2010 A wnt oscillator model for somitogenesis *Biophys. J.* **98** 943–50
- [120] Petrie C, Anderson C, Maekawa C, Maekawa T and Transtrum M K 2021 The supremum principle selects simple, transferable models (arXiv:2109.10449)
- [121] Knight-Schrijver V, Chelliah V, Cucurull-Sanchez L and Le Novère N 2016 The promises of quantitative systems pharmacology modelling for drug development *Comput. Struct. Biotechnol. J.* **14** 363–70
- [122] Transtrum M K 2020 Using simple models to understand complex biological processes *Food and Drug Administration Seminar*
- [123] Peterson M and Riggs M 2015 FDA advisory meeting clinical pharmacology review utilizes a quantitative systems pharmacology (QSP) model: a watershed moment? *CPT: Pharmacomet. Syst. Pharmacol.* **4** 189–92
- [124] Curtarolo S *et al* 2012 AFLOW: an automatic framework for high-throughput materials discovery *Comput. Mater. Sci.* **58** 218–26
- [125] Saal J E, Kirklin S, Aykol M, Meredig B and Wolverton C 2013 Materials design and discovery with high-throughput density functional theory: the open quantum materials database (OQMD) *JOM* **65** 1501–9
- [126] Tadmor E B, Elliott R S, Sethna J P, Miller R E and Becker C A 2011 The potential of atomistic simulations and the knowledgebase of interatomic models *JOM* **63** 17
- [127] Holzinger A 2018 From machine learning to explainable AI *2018 World Symp. on Digital Intelligence for Systems and Machines (DISA)* (IEEE) pp 55–66
- [128] Hagrass H 2018 Toward human-understandable, explainable AI *Computer* **51** 28–36
- [129] Pawar S, San O, Aksoylu B, Rasheed A and Kvamsdal T 2021 Physics guided machine learning using simplified theories *Phys. Fluids* **33** 011701
- [130] Willard J, Jia X, Xu S, Steinbach M and Kumar V 2020 Integrating scientific knowledge with machine learning for engineering and environmental systems (arXiv:2003.04919)
- [131] Rai R and Sahu C K 2020 Driven by data or derived through physics? A review of hybrid physics guided machine learning techniques with cyber-physical system (CPS) focus *IEEE Access* **8** 71050–73
- [132] Andersson G *et al* 2005 Causes of the 2003 major grid blackouts in North America and Europe and recommended means to improve system dynamic performance *IEEE Trans. Power Syst.* **20** 1922–8
- [133] Laughlin R B and Pines D 2000 From the cover: the theory of everything *Proc. Natl Acad. Sci. USA* **97** 28
- [134] Crutchfield J P 2014 The dreams of theory *Wiley Interdiscip. Rev.-Comput. Stat.* **6** 75–79