scientific reports

Check for updates

OPEN

Predictive power of non-identifiable models

Frederic Grabowski, Paweł Nałęcz-Jawecki & Tomasz Lipniacki✉

Resolving practical non-identifiability of computational models typically requires either additional data or non-algorithmic model reduction, which frequently results in models containing parameters lacking direct interpretation. Here, instead of reducing models, we explore an alternative, Bayesian approach, and quantify the predictive power of non-identifiable models. We considered an example biochemical signalling cascade model as well as its mechanical analogue. For these models, we demonstrated that by measuring a single variable in response to a properly chosen stimulation protocol, the dimensionality of the parameter space is reduced, which allows for predicting the measured variable's trajectory in response to different stimulation protocols even if all model parameters remain unidentified. Moreover, one can predict how such a trajectory will transform in the case of a multiplicative change of an arbitrary model parameter. Successive measurements of remaining variables further reduce the dimensionality of the parameter space and enable new predictions. We analysed potential pitfalls of the proposed approach that can arise when the investigated model is oversimplified, incorrect, or when the training protocol is inadequate. The main advantage of the suggested iterative approach is that the predictive power of the model can be assessed and practically utilised at each step.

Non-identifiable or sloppy models are ubiquitous in systems biology[1,2]. Non-identifiability, in practical terms, means that parameters obtained through model fitting cannot be trusted: either because arbitrarily large changes of some parameters can be fully compensated by changes of other parameters (structural non-identifiability) or because plausible parameters may differ widely due to noise in measurements (practical non-identifiability). A closely related property is sloppiness. In sloppy models, some combinations of parameters define stiff directions, along which model responses change rapidly, whereas the remaining combinations define sloppy directions, along which model responses change insignificantly. It is important to note that non-identifiability or slopiness is a property of a model together with a dataset, not just the model itself. Additional data can make the model (more) identifiable[3] and can reduce the number of sloppy directions.

Non-identifiability and sloppiness have a lot in common, however the first is perceived as a model disadvantage that should be resolved, whereas the second is perceived as a ubiquitous property of complex models we have to live with[2]. Typically, structural non-identifiability is relatively easy to resolve by nondimensionalisation and reparametrization. Practical non-identifiability is a harder problem, usually requiring additional experiments (which can be expensive or time-consuming), or manual model reduction[4,5]. An identifiable model that reflects underlying biochemistry can be seen as the ultimate goal of systems biology. However, model reduction, performed to reach identifiability, may lead to composite parameters that lack straightforward biochemical interpretation[6–8]. Additionally, a reduced model may not be able to accommodate new data containing measurements of previously omitted variables. An alternative approach to non-identifiable, sloppy models is to identify stiff directions, and analyse system behaviour along these directions, which is an effective way of reducing the dimensionality of the investigated parameter space, without changing the model[9].

In this study, we focus on practically non-identifiable models, but instead of reducing them to reach identifiability, we train them on limited datasets and investigate their predictive power after such training. It was observed earlier by Brown et al.[10], Cedersund[11], van Mourik et al.[12], Monsalve-Bravo et al.[9] and others that sloppy, non-identifiable models may lead to informative predictions; nevertheless the standard solution is to make models identifiable before predictions are made[3], as reviewed in Refs.[4,13]. In this paper, we analysed the predictive power of non-identifiable models in more detail in the context of systems biology and investigated how successive inclusion of model variables in training influences predictions[14]. We view modelling as the following iterative procedure: perform experiment → train model on gathered data → assess its predictive power, if required perform additional experiments → train model again, narrowing down its plausible parameter space and thus increasing its predictive power. This approach stems from the fact that in (systems) biology additional

Institute of Fundamental Technological Research, Polish Academy of Sciences, Warsaw, Poland. ✉email: tlipnia@ippt.pan.pl

measurements are usually relatively expensive and/or time-consuming, so researchers try to constrain the model by performing a minimal necessary set of experiments[15,16]. This procedure increases the number of stiff directions and ultimately allows one to predict the investigated variables with high accuracy. To illustrate the procedure outlined above we considered a biochemical signalling cascade with a negative feedback loop—a motif ubiquitous in regulatory pathways (e.g. MAPK[17,18], NF-κB[19,20], p53[21,22]). As an auxiliary example, we analysed a simple damped mechanical oscillator consisting of three masses with one of them subject to forcing.

We investigated the trained model's ability to predict the system responses to different temporal stimulation protocols as well as responses that arise when one of the parameters is multiplied or divided by a factor of 10. The latter corresponds to perturbations that can be introduced by inhibitors or mutations causing protein hyperactivity, while the change of stimulation protocol can be interpreted as a modification of the therapy schedule. In the last decade, investigation of responses to complex temporal stimulations was enabled by optogenetic receptors which allow arbitrarily complex stimulation protocols[23,24]. We demonstrated that training a relaxed model, with three hypothetical negative feedback loops, leads to equally good predictions as training the nominal model, and correctly suggests the location of the nominal feedback loop. In contrast, training oversimplified or incorrect models leads to inaccurate predictions with (misleadingly) narrow prediction bands. This suggests the use of a more relaxed model structure when the exact network wiring is unknown or uncertain.

Model training relies on the exploration of the space of plausible parameters, i.e. parameters giving predictions within an expected error bound of the measurements. There are several methods to explore the parameter space (e.g. Fisher information matrix[16,25], profile likelihood[3,4], Bayesian methods[26–28], approximate Bayesian computation[29], nested sampling[30]). We chose the Bayesian approach and sampled the parameter space by performing Markov Chain Monte Carlo (MCMC) simulations using the Metropolis–Hastings algorithm[31]. This method allows us to start from broad, non-informative priors and explore the posterior parameter distributions confined by smaller or larger datasets, and consequently obtain straightforward probabilistic predictions. To avoid over-focusing on the details of the Bayesian approach, we will refer to posterior parameter samples as the set of plausible parameters, and to the process of running the MCMC simulation and obtaining the samples as training.

Results

Sequential training of the signalling cascade model.

We considered a four-step signalling cascade with negative feedback(s) resembling the MAPK cascade (RAS → RAF → MEK → ERK), outlined in Fig. 1a, with corresponding equations given in Fig. 1b. The cascade is activated by a time-dependent signal $S(t)$, which depends on the stimulation protocol chosen. Nominal model parameters (used for generating training data) and
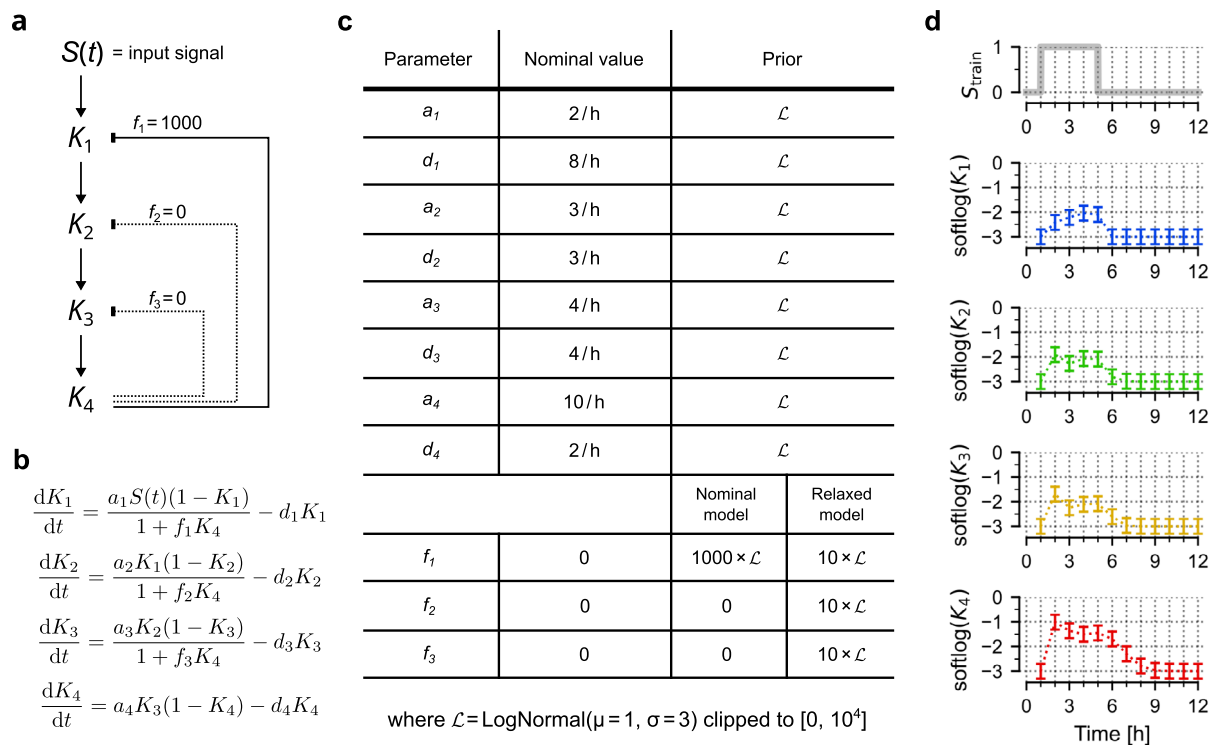
Figure 1. Signalling cascade model. (**a**) Scheme of the 4-level biochemical signalling cascade model activated by a time-dependent signal $S(t)$. Strengths of feedbacks f_2 and f_3 are set to zero in the nominal model but are allowed to assume positive values in the 'relaxed' model. (**b**) Model equations. (**c**) Model parameters used for generating training data and prior distributions used for training. (**d**) Simulated trajectories of K_1, K_2, K_3, and K_4 in response to S_{train} used for model training; softlog$(x) = \log(0.001 + x)$. Error bars show measurement errors assumed for the generation of training data.

assumed lognormal parameter priors used during model training are given in Fig. 1c. We note that, by default, the strengths of feedbacks to the second and third step of the cascade are zero ($f_2 = f_3 = 0$)—later we investigate a 'relaxed' model in which f_2 and f_3 are allowed to be positive. The training dataset was generated by simulating trajectories of the four model variables, K_1, K_2, K_3, and K_4 using an 'on–off' stimulation protocol and randomly perturbing them in order to mimic experimental errors (Fig. 1d). The error bars are plotted at measurement time points (every hour) and show the magnitude of the assumed experimental errors. We used three such generated measurement replicates for training (see "Methods" for details). The training dataset was used to train two models: the 'nominal' model for which we assume $f_2 = f_3 = 0$, and a 'relaxed' model (in the next section) in which we allow f_2 and f_3 to assume positive values, also with a lognormal prior (Fig. 1c).

First, we trained the nominal model using only the trajectory of the last variable, K_4. We observe that the resulting model accurately predicts the trajectory of K_4 in response to two very different stimulation protocols (Fig. 2a,b), but fails to predict the trajectories of any of the remaining variables. For these three variables, the 80% prediction bands are very broad, implying that the trained model contains little information about these variables. Next, the training dataset was expanded to additionally include the trajectory of K_2, and we observe that also the trajectory of this variable can be reliably predicted (Fig. 2c). Finally, the complete training dataset containing trajectories of all four variables resulted in a "well trained" model, which predicts the trajectories of all variables (Fig. 2d).

By investigating the space of plausible parameters (i.e., posterior samples) we may notice that even for the well-trained model the marginal probability distributions show that all 9 parameters may vary by about an order of magnitude, and when training is performed only on K_4 (or K_2 and K_4) parameters may vary by about two orders of magnitude (Fig. 2e). This shows that the model trained only on the last variable of the cascade, K_4, can accurately predict this variable's trajectories in response to different stimulation protocols, even though all model parameters remain uncertain, suggesting sloppiness of the trained model[10]. To better understand this feature we numerically analysed the dimensionality of the plausible parameter sets (see "Methods" section for details) obtained after training the model on (i) K_4, (ii) K_2 and K_4, and (iii) all four variables (Fig. 2f). Since we are interested in multiplicative changes of parameters, we performed a principal component analysis (PCA) on sets of (natural) logarithms of plausible parameters[9]. Exponents of the square roots of the principal values (variances) λ_i can be interpreted as principal multiplicative deviations of parameter values, $\delta_i = \exp(\sqrt{\lambda_i})$. The prior parameter distributions are LogNormal($\mu = 1$, $\sigma = 3$) clipped to $[0, 10^4]$, which gives $\delta_{i,\text{prior}} \approx \exp(3) \approx 20$ (for all i), implying that "on average" all prior parameter values can vary 20-fold from the median. After training on all variables, the largest principal multiplicative deviation is $\delta_i \approx 12$, however, the 4 smallest δ_i are much lower, below 1.5 (note that $\delta_i = 1$ implies no variation along the i-th principal component), indicating that these directions are stiff. This implies that the dimensionality of the plausible parameter space was effectively reduced by 4 due to training. When training is performed only on K_4, dimensionality is reduced by 1 (only the smallest δ_i is below 1.5), and when training uses K_2 and K_4, dimensionality is reduced by 2 (the two smallest δ_i are below 1.5). Therefore, in agreement with intuition, model training reduces the dimensionality of the plausible parameter space. In our example, measurement of n variables reduces dimensionality by n. Importantly, even reduction from 9 to 8 dimensions allows for reliable predictions of a single variable of interest that was measured to train the model.

The principal components (eigenvectors) can be used to gain additional insight. Each stiff direction defines a combination of parameters that is well-constrained by data used for training (Fig. S1a). The single stiff direction that emerges after training on K_4 has four positive coordinates corresponding to activation parameters (a_1–a_4) and five negative coordinates corresponding to deactivation parameters (d_1–d_4) and the negative feedback (f_1). The same pattern appears when training includes two, or all four variables; in the emerging stiff directions, coordinates corresponding to activation and deactivation have opposite signs. Stiff directions span a stiff subspace, and this subspace can be more intuitively visualised when a linear combination of stiff vectors minimising the number of dominant coordinates is chosen (Fig. S1b). One can now readily see that even after training on all four model variables, only the ratios of activation/deactivation parameters can be determined, not these parameters themselves.

Finally, we investigated whether the trained model is able to predict responses to test protocols that arise after a multiplicative change of a given model parameter. Such a change is equivalent to a system perturbation that can be introduced experimentally by knockdown or overexpression of a reacting protein. To test this, after training the model on K_4, we modified each set of plausible parameters by multiplying or dividing a given parameter by a factor of 10. As shown in Fig. 3, for each of the 9 parameters such a modification allows to correctly predict, with relatively narrow prediction bands, the response of the nominal model with the corresponding parameter change.

Training a relaxed model.

Frequently, experiments are performed not only to constrain the parameters of a known model but to deduce the structure of the model itself. To explore this scenario, in Fig. 4 we introduced a 'relaxed' model, in which we allowed strengths of the three negative feedbacks shown in Fig. 1a to assume positive values during training (assuming prior lognormal distributions, as shown in Fig. 1c). This increases the number of free parameters from 9 to 11. Again, this relaxed model was trained on in silico-generated data, based on the nominal model with one feedback. Regardless of whether only K_4 or all variables are included in the training dataset, the relaxed model performs equally well as the nominal model on test protocols (compare Fig. 4a versus Figs. 2a, 4b versus Fig. 2d). We note that the marginal distributions of the three feedback parameters remain broad (Fig. 4c). This means that, even without knowing strengths of individual feedbacks, the model can give reasonable predictions. Nevertheless, these distributions indicate that the (nominal) feedback f_1 is dominant.
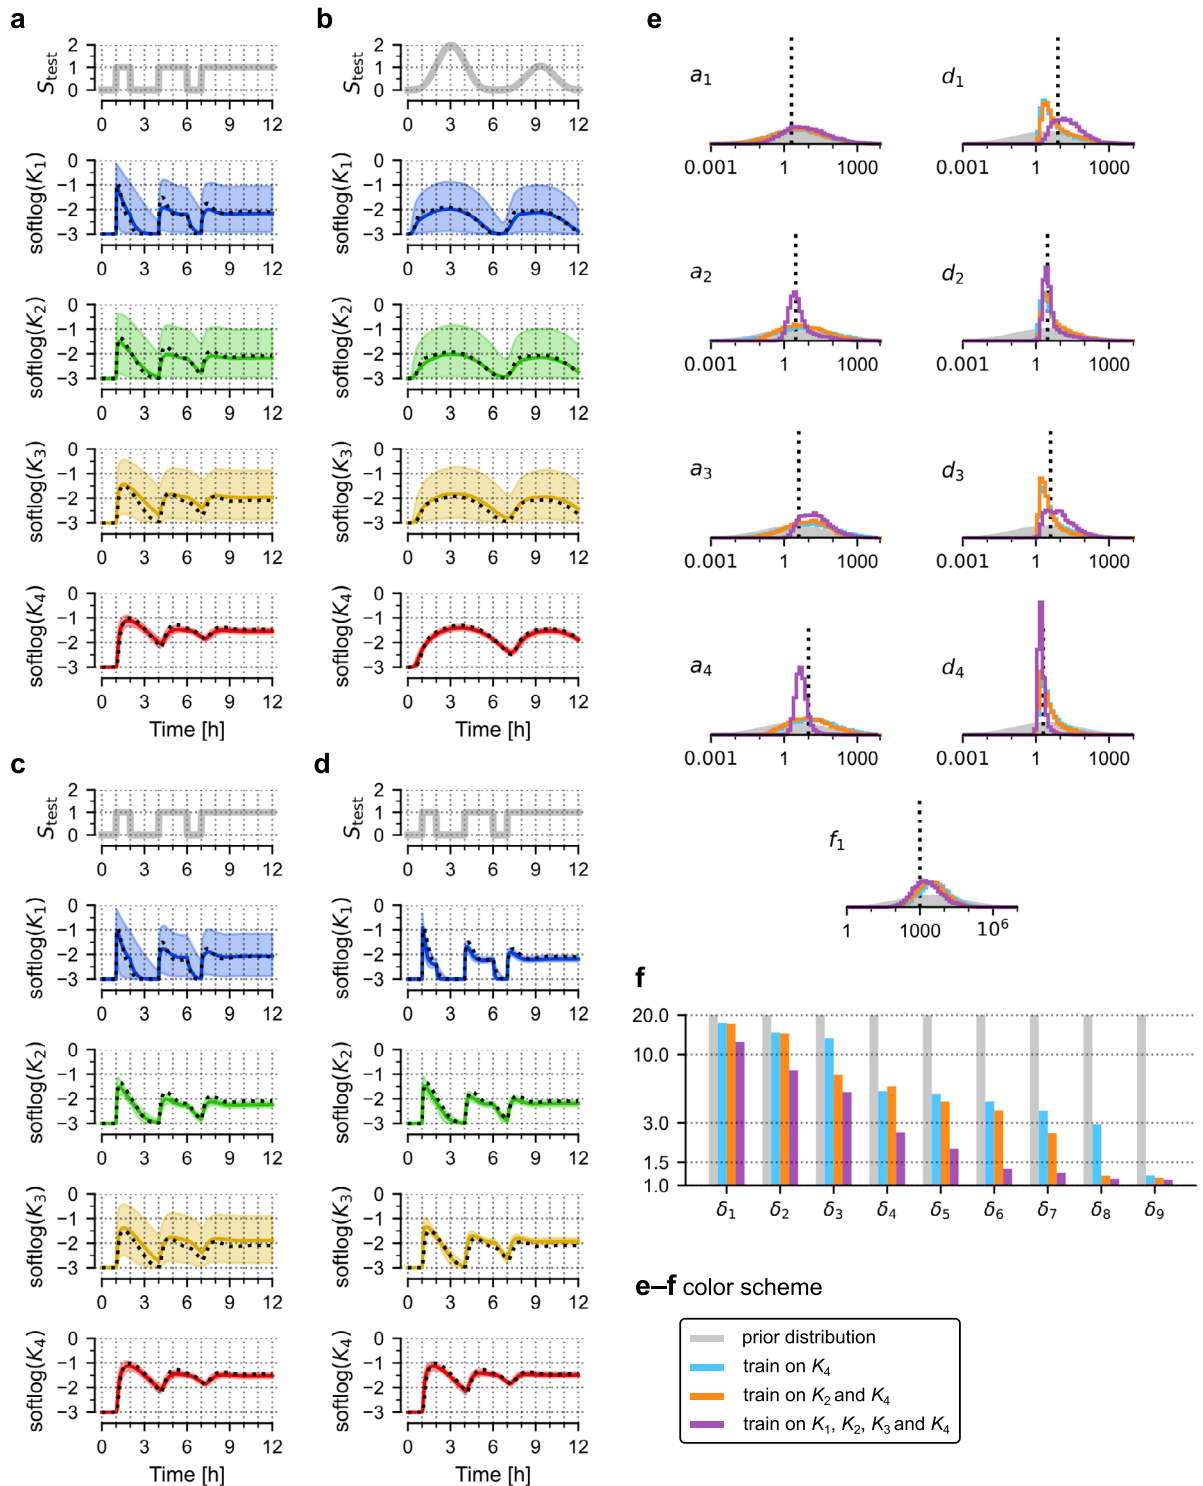
Figure 2. Training the nominal model. (**a,b**) Prediction of model responses to two test signals, S_{test}, after training on the trajectory of K_4 (generated in response to the training protocol, S_{train}, shown in Fig. 1d). Black dotted lines show trajectories of the nominal model, coloured lines show (point-wise) medians of predictions, contours show 80% prediction bands; softlog(x) = log(0.001 + x). (**c**) Prediction of model responses after training on trajectories K_2 and K_4. (**d**) Prediction of model responses after training on trajectories of all four model variables, K_1, K_2, K_3 and K_4. (**e**) Histograms obtained from 10,000 samples generated from the posterior distribution of model parameters after training on K_4 (blue); K_2 and K_4 (orange); K_1, K_2, K_3 and K_4 (purple). Grey contours show prior distributions. (**f**) Dimensionality reduction of the parameter space by successive inclusion of model variable trajectories into the training set. Bar heights show principal multiplicative deviations $\delta_i = \exp(\sqrt{\lambda_i})$, where λ_i are principal values (variances) of log-parameters.
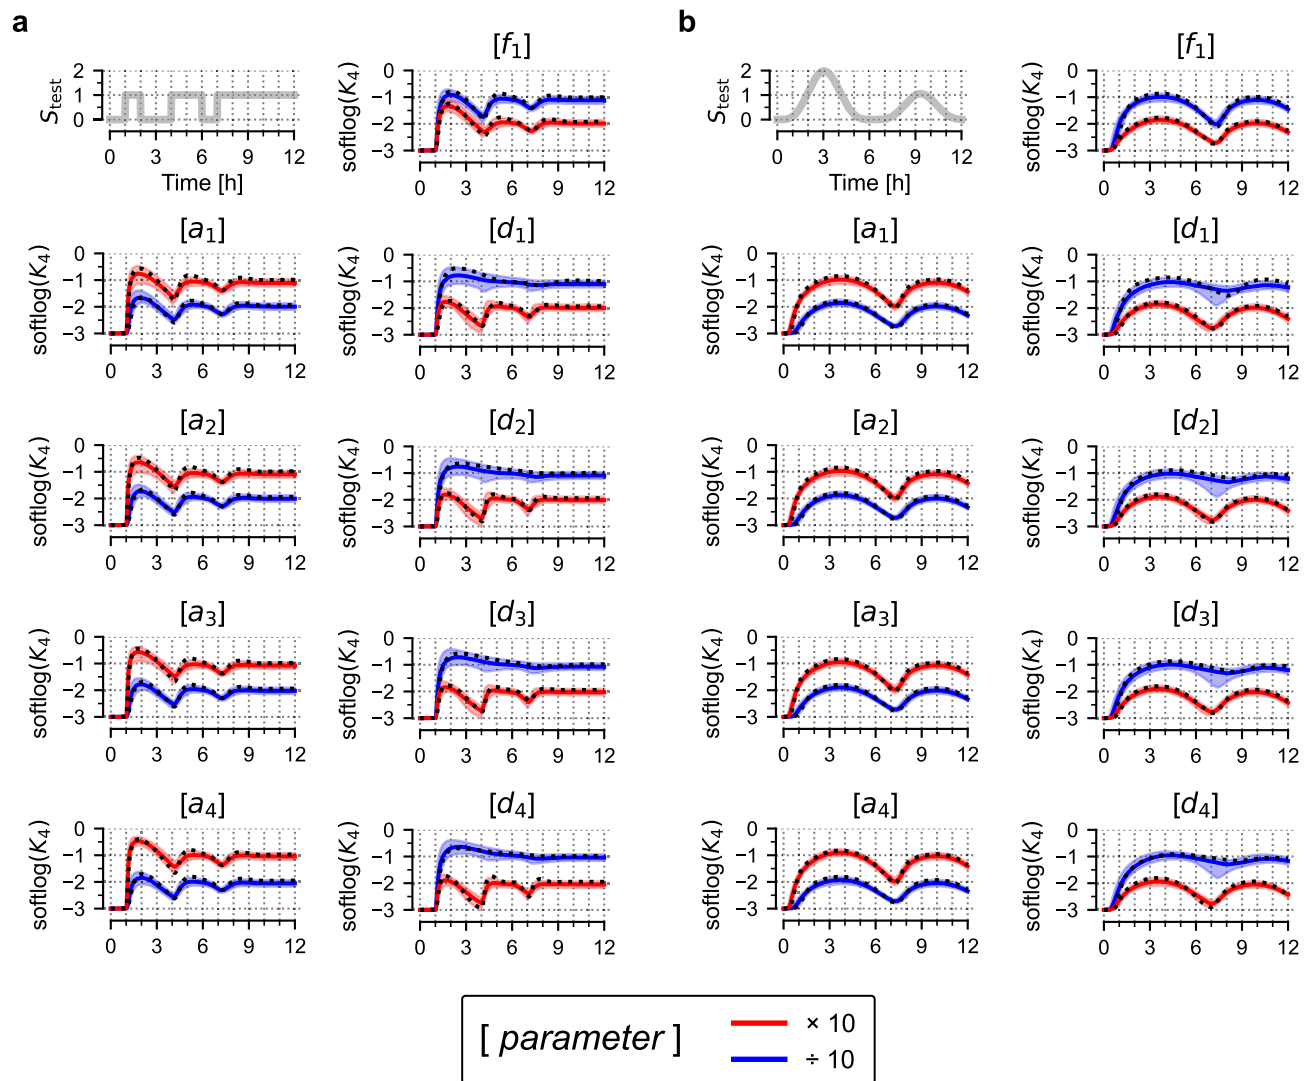
Figure 3. Predicted responses after changing model parameters. (**a,b**) Prediction of model responses to two test signals, S_{test} (shown in each panel top left), after training on the trajectory of K_4 (generated in response to the training protocol, S_{train}, shown in Fig. 1d). In each chart, one of the parameters (shown in square brackets) was multiplied (red) or divided (blue) by a factor of 10. Black dotted lines show trajectories of the nominal model, coloured lines show (point-wise) medians of predictions, contours show 80% prediction bands; $softlog(x) = log(0.001 + x)$.

The above example suggests that if there exists uncertainty about model structure, it is advisable to train a more relaxed model, which can still result in satisfactory predictions, and eventually, after gathering more data, can provide a hint about true model structure.

Training a simplified model. Given insufficient knowledge about the cascade structure and a small dataset, it might be tempting to train a simplified model. Here, we give an example that showcases potential pitfalls of this approach. As previously, the training set was a perturbed trajectory of K_4 generated by the nominal model—however, we attempted to train a model that contains only two signalling steps, in which K_1 directly influences K_4 (formally in the equation for K_4 we replace K_3 by K_1, and equations for K_2 and K_3 are omitted). Training the simplified model on the 'on–off' protocol shows some discrepancy between the original and simplified model trajectories (Fig. 5a, note the logarithmic scale of the plot). When the trained model is verified on the test stimulation protocol, the observed discrepancy is more visible, meaning that the simplified model is not able to make accurate predictions (Fig. 5b). Crucially, the 80% prediction band is relatively narrow, similar to that obtained for the correct model (Fig. 5b versus Fig. 2a). Thus the trained simplified model is "sure" (even if marginal parameter distributions of 3 out of 5 model parameters remain very broad—Fig. 5c) about its inaccurate predictions.

In certain cases where approximate predictions are still satisfactory, the simplified model can be a reasonable solution. However, as demonstrated above, one should be very careful when interpreting parameters and making predictions, since some conclusions might be an artefact of using a wrong model to compute posterior probabilities. Finally, in contrast to the relaxed model, the simplified model cannot be improved by additional data.
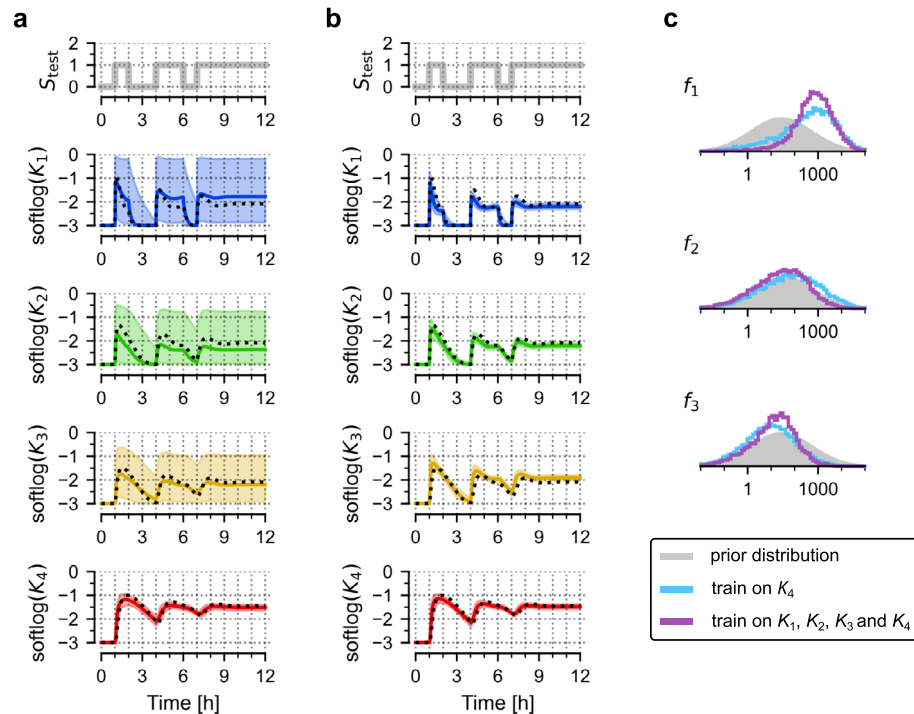
Figure 4. Training a relaxed model. (**a**) Prediction of model responses to test signal, S_{test}, after training on the trajectory of K_4 (generated in response to the training protocol, S_{train}, shown in Fig. 1d). Black dotted lines show trajectories of the nominal model, coloured lines show (point-wise) medians of predictions, contours show 80% prediction bands; softlog$(x) = \log(0.001 + x)$. (**b**) Prediction of model responses after training on trajectories of all four model variables. (**c**) Histograms were obtained from 10,000 samples generated from the posterior distribution of model parameters after training on K_4 (blue) or all model variables (purple). Grey contours show prior distributions.

Analysis of potential pitfalls. In Fig. S2 we investigate training based on an inadequate stimulation protocol. The protocol's 'on' phase lasts only 1 h (compared to 4 h in Fig. 2), and as a result, the model predictions have a broad 80% prediction band. The prediction band for K_4 is broad when training includes K_4 only, and remains broad even when training includes all four model variables. Higher uncertainty of trained model trajectories is associated with larger principal multiplicative deviations δ_i, compared to values obtained for training shown in Fig. 1d. This example highlights the importance of designing a good stimulation protocol for training. If our goal is to obtain accurate predictions of K_4 only, training on K_4 with the help of an appropriate protocol (Fig. 2a) gives much better results than training on all variables employing a suboptimal protocol (Fig. S2c). This pitfall is relatively simple to resolve—one has to adjust the training protocol so that it allows capturing characteristic temporal responses of the system. Whether the 1-h long 'on' phase is sufficient or not depends on the nominal model parameters.

In Fig. S3 we show an example where the prior model is incorrect, i.e. has no negative feedback at all. We obtain discrepancies between measurement and model, which could be overlooked during training if, for example, the measurement points were not distributed densely enough. From the trained wrong model we also obtain wrong predictions with very narrow prediction bands, as in the case of the simplified model discussed earlier. The key message from Fig. 5 and Fig. S3 is that a simplified or wrong model can make inaccurate or wrong predictions with narrow prediction bands, and thus the width of prediction bands contains no information about the accuracy of the model. Narrow prediction bands imply that obtained constraints on parameters suffice for unique predictions, but these predictions are correct only under the condition that the model structure is also correct.

Mechanical analogue. We performed a similar analysis of a (distant) mechanical analogue of the signalling pathway, a damped harmonic oscillator consisting of three masses connected by two springs (Fig. 6). The system is parametrized by 8 free parameters. For training, we used a protocol in which one of the masses is pushed with constant force for 10 s, which results in a displacement of all masses. Next, we used the trajectory of the opposite mass (or of all masses) to train the model. In the test protocol, a sinusoidal driving force was applied to the same mass as during training. Similarly to the signalling cascade model, we found that training on a single mass trajectory suffices to predict its trajectory (for a limited time) in response to a very different forcing, and training on all three masses allows for very accurate predictions of all trajectories, despite the fact that marginal distributions of model parameters remain broad.
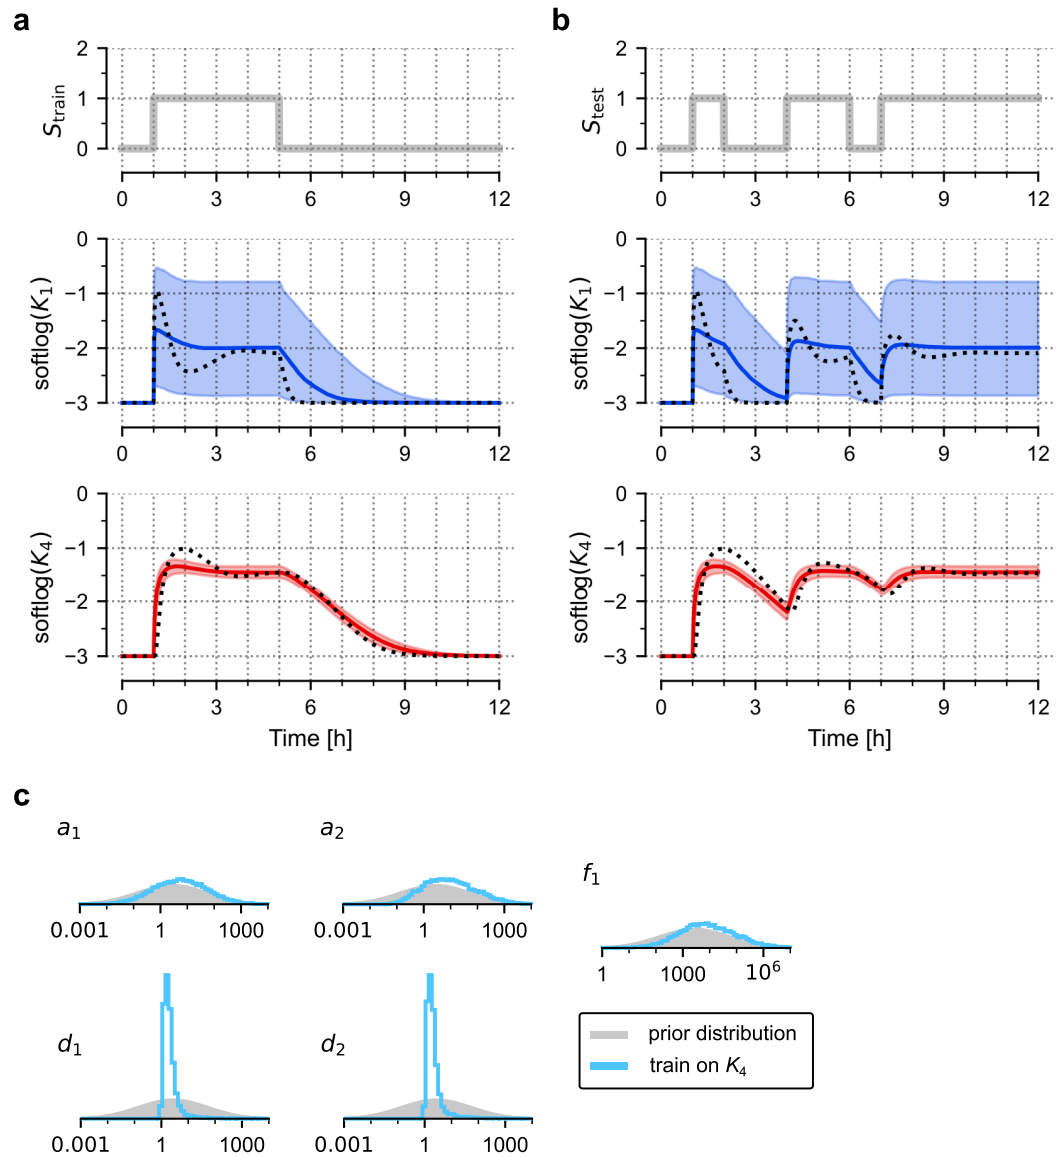
Figure 5. Training a simplified model. (**a**) Training a simplified (2-step) model on the nominal (4-step) model trajectory of K_4. Only variables present in the simplified model are shown. Black dotted lines show trajectories of the nominal model, coloured lines show (point-wise) medians of the simplified model, contours show 80% prediction bands; $\mathrm{softlog}(x) = \log(0.001 + x)$. (**b**) Prediction of the trained simplified model responses to the test signal, S_{test}. Same graphical convention as in (**a**). (**c**) Histograms were obtained from 10,000 samples generated from the posterior distribution of simplified model parameters. Grey contours show prior distributions.

Discussion

The majority of systems biology models, including mechanistic models of intensely studied regulatory networks, are non-identifiable (or sloppy). Current techniques and algorithms allow practitioners to automatically resolve structural non-identifiability problems, however, resolving practical non-identifiability is harder, and frequently requires additional data or substantial model reduction. We see the model reduction problem as follows: regulatory networks have built-in signalling cascades that transmit signals between key nodes. In many cases, researchers are aware of components (typically kinases) of these signalling cascades, but measuring the kinetics of all these components is not feasible. It is thus tempting to skip these intermediate steps, but this solution is problematic since these intermediate steps introduce time delays that can be important for whole-network behaviour, such as oscillations. In the case of the discussed cascade $S \rightarrow K_1 \rightarrow K_2 \rightarrow K_3 \rightarrow K_4$ triggered in $t = 0$ from state $K_1 = K_2 = K_3 = K_4 = 0$ by the step signal $S(t)$, the first, second and third time derivatives of K_4 (at $t = 0$) are zero. When the cascade is reduced to $S \rightarrow K_1 \rightarrow K_4$, the second derivative of K_4 becomes positive, which noticeably changes the system's response, as shown in Fig. 5. To maintain observed dynamics (in particular oscillations) one can introduce nonlinearities, but as a consequence, the resulting model misses some components and has nonlinear terms not supported by underlying biochemistry[32,33].
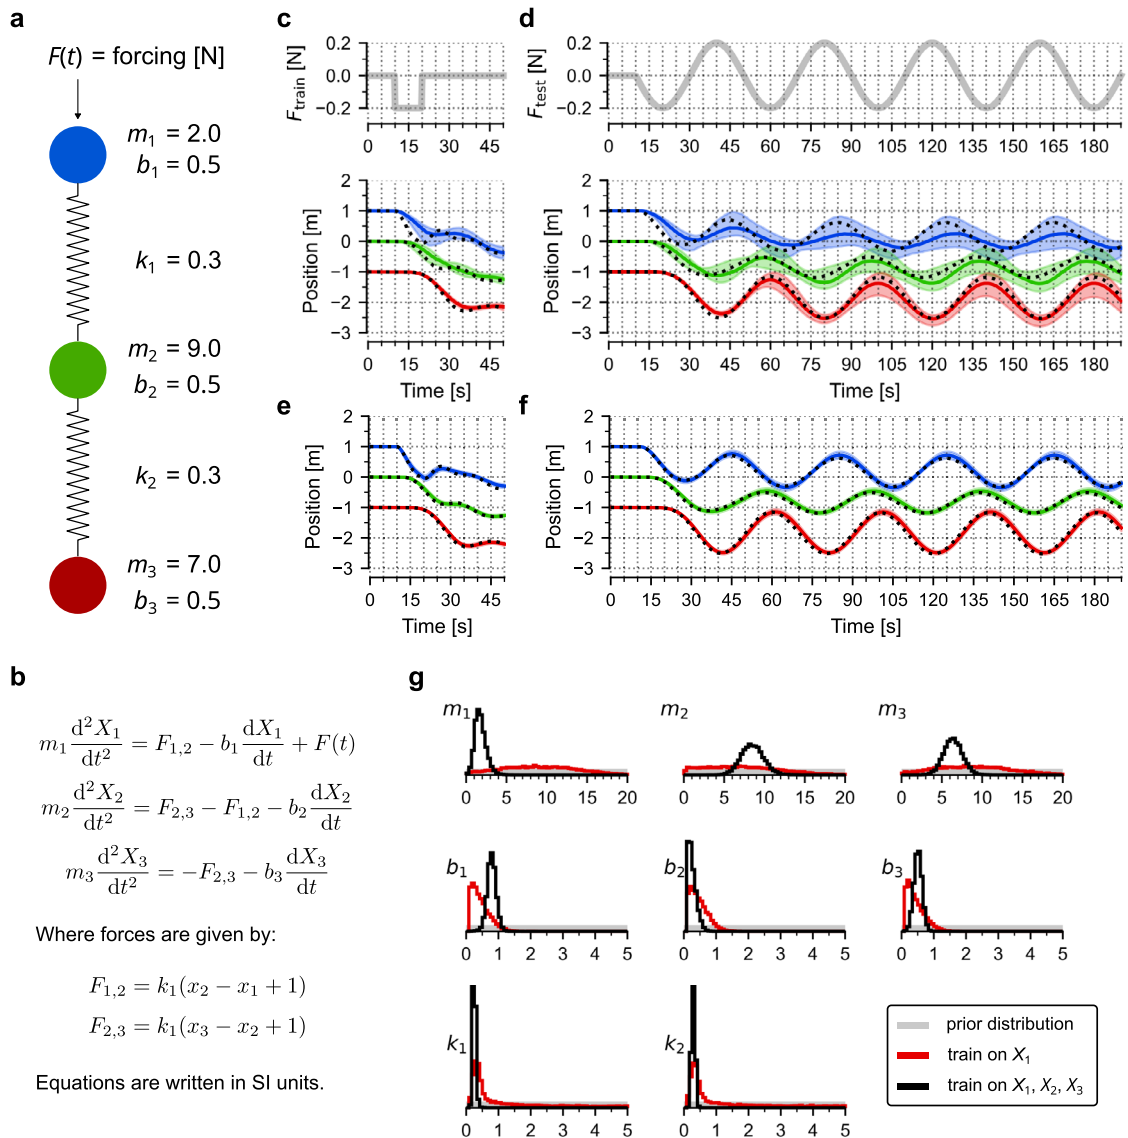
Figure 6. Three-body harmonic oscillator. (**a**) Scheme of the model. All eight parameters (masses, m_i; damping coefficients b_i; spring constants k_i) to be estimated are given together with their nominal values. Before the time-dependent force $F(t)$ is applied, the system remains in equilibrium. (**b**) Model equations. (**c,d**) Training (**c**) and prediction (**d**), with only the trajectory of the red mass included in the training. As in the signalling cascade model, for training, we use three independently generated measurement replicates with Gaussian noise (sigma $\sigma_{\text{error}} = 0.3$) added to the deterministic trajectory (black dotted line). Measurement points are at full seconds. Coloured lines show (point-wise) medians of predictions, contours show 80% prediction bands. (**e,f**) Training (**e**) and prediction (**f**), with trajectories of all masses included in the training, same convention as in (**c,d**). (**g**) Histograms were obtained from 10,000 samples generated from the posterior distribution of model parameters after training on red mass trajectory (red line) and after training on all masses trajectories (black line). Grey contours show (uniform) prior distributions.

In response to the outlined problems with model reduction, we studied an alternative approach based on training the full mechanistic model on the available dataset and assessing its predictive power[2]. We showed that a very limited dataset, insufficient to constrain any model parameters, can still allow for satisfactory predictions[10,11]. In particular, we demonstrated that by training the signalling cascade model on the trajectory of the last component of the cascade (obtained in response to some properly chosen training protocol), one can unambiguously predict its trajectory in response to other protocols, as well as parameter changes. This prediction is possible despite the fact that trajectories of all upstream variables remain unknown.

Investigation of the plausible parameter set allows for additional insight. In the considered example we observe that when training is performed on trajectories of n variables, the 'effective' dimensionality of the plausible parameter set is reduced by n, and this reduction suffices to predict trajectories of the measured variables in response to other protocols. As discussed by Machta et al.[34], the compression of the parameter space identifies stiff directions corresponding to composite parameters controlling observed variables. By picking a suitable basis for

8

the stiff space, we confirmed our intuition that ratios of activation/deactivation parameters govern the dynamics of the cascade. Since stiff directions usually do not align with model parameters (as it is in our case), Gutenkunst et al. argued that in sloppy models, measuring system dynamics is an ineffective way of deriving parameter values, and conversely, measuring parameter values is an ineffective way of predicting system dynamics[2]. In fact, even after training on trajectories of all model variables, most model parameters remained unidentified.

Model training can be generally understood as exploring the parameter space, in order to find plausible parameters, i.e. parameters leading to model trajectories in agreement with available data. In this context, many exploration methods are available. Locally, information about the uncertainty of parameters can be calculated using Fisher's information matrix. For models with asymmetric but still unimodal loss landscapes, the (semi-local) profile likelihood method can be used. For arbitrary loss landscapes, a Bayesian approach together with MCMC exploration is convenient, as

1. it allows for the inclusion of prior knowledge; if no such knowledge exists, training can start from some broad uninformative prior,
2. it allows for the analysis of how much a model has learnt by including a given dataset; this can be expressed by the reduction of the 'effective' dimensionality of the plausible parameter space, or entropy of the distribution of plausible parameters,
3. it allows us to assess the predictive power of the trained model, by sampling the posterior parameter distribution and computing prediction bands for variable trajectories.

The main disadvantage of MCMC compared to local (or semi-local) methods is that it can become computationally expensive, especially for high-dimensional parameter spaces[28]. The computational cost is typically lower if exploration starts from a relatively narrow (informative) prior, or a more complete dataset is available, which results in a narrower posterior distribution[35]. A promising method of accelerating MCMC exploration is based on reducing expensive model evaluations by local approximation of the likelihood[36,37].

In summary, the advantage of the studied approach to non-identifiability is twofold.

1. It does not require pre-emptive model reduction to avoid unmeasured variables in order to obtain predictions. Certainty of the predictions obtained from the full model can be determined by predictive bands, and if these predictions are insufficiently narrow, the model can be further trained using additional measurements of the variables of interest. Notably, these predictions are more reliable than predictions obtained from simplified models. While predictions from the full model may have wide predictive bands, the predictions from simplified models can be misleading (inaccurate with narrow bands) due to introduced simplifications.
2. If at a later time, measurements of previously unobserved variables are performed, no modification of the model structure is necessary. In contrast, reduced models may not be able to accommodate new data, simply because some components were omitted during reduction.

We hope that this approach will allow for easier collaboration, as it allows researchers to jointly work on the same model, improving it by performing new experiments[2].

Identifiable models or theories stem from physics, where they showed their superiority over other approaches. Whole branches of physics are based on a small number of fundamental theories. In systems biology, the situation is different: there are a lot of not-so-important models, typically having a large number of parameters that can neither be (feasibly) derived from more fundamental theories nor effectively measured. Currently, a very intensively developed approach to complex systems is based on neural networks. In this approach, one gets little to no insight into how the system works but can get very good predictions. In this sense, Bayesian methods are a middle ground between classical, identifiable models and black box models such as neural networks. The Bayesian approach studied here works well in our case because we have a good prior on the model structure, even if the prior on model parameters is broad and uninformative. This approach can be impractical for very complex systems for which neural networks are typically used, but is able to both give useful predictions and obtain some insight into the underlying regulatory system. Even if the explanatory power of the trained non-identifiable model is limited by the uncertainty of its parameters, its predictive power can make it useful (see also discussion in Refs.[2,38]).

We think that the discussed approach, based on model training on some selected variables and giving only partial insight into the system's behaviour, may reflect true uncertainties in underlying biology[39,40]. Considering a regulatory network one can identify core nodes. Because of their importance, the corresponding variables are frequently measured, and for the same reason one can expect relatively good replicability of these measurements. The intermediate pathway components are less tightly constrained and may exhibit higher variability[41]. Daniels et al. argued that cell-to-cell variability in protein levels might be beneficial for robustness and evolvability[39]. Variability of intermediate proteins implies that (for very practical reasons) these variables are not (frequently) measured; or even if they are measured, results are (understandably) not reported due to insufficient replicability. One can thus hypothesise, that in the case of intensively investigated regulatory systems, the remaining uncertainty of model parameters and of trajectories of intermediate components may reflect their true variability, at least at the single cell level. One may hope that quantification of the variability of intermediate variables, and correlation between them will help to elucidate existing compensatory properties ensuring robustness of the whole regulatory network.

Methods

Computational model simulations.   Each model is represented by a set of ODEs and an input signal $S(t)$, which are solved using an explicit Runge–Kutta (4,5) formula developed by Dormand-Prince. Model simulations are used first to generate in silico measurements, and then to estimate parameter distributions based on these measurements. A measurement trajectory is evaluated at full hours (or at full seconds in the case of the mechanistic analogue model shown in Fig. 6) by adding to the ODE solution Gaussian noise (with sigma $\sigma_{error} = 0.3$) that emulates measurement error. Parameter estimation (model training) is based on 3 independent measurement trajectories.

In the case of biochemical cascade models, ODE trajectory $x(t)$ is log-transformed to $y(t) = \log_{10}(0.001 + x)$ before adding measurement noise.

Parameter estimation based on MCMC simulations.   For each of the considered models we assume some prior parameter distribution $P(\theta)$ and use an MCMC algorithm to draw samples from the posterior parameter distribution $P(\theta \mid M)$, given in silico generated measurements M.

To estimate $P(\theta \mid M)$, we first calculate the likelihood of θ given a single measurement $L_{k,t,r}(\theta) = P(m_{k,t,r} \mid \theta)$ (where k is the measured variable, t is the time of measurement, and $r = 1, 2, 3$ is the measurement replicate) by simulating variable trajectories $y_k(t)$ for the parameter set θ. Then

$$L_{k,t,r}(\theta) = \frac{1}{\sigma_{error}\sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma_{error}^2}(y_k(t) - m_{k,t,r})^2\right),$$

where σ_{error} is the measurement error (assumed 0.3 for all our models). The total likelihood given all measurements of variable k is thus:

$$L_k(\theta) = \prod_{t,r} L_{k,t,r}(\theta),$$

and the posterior distribution $P(\theta \mid M)$ is proportional to

$$f(\theta) = P(\theta)P(M|\theta) = P(\theta)\prod_k L_k(\theta),$$

where the product is over variables k selected for model training. Then $f(\theta)$ is used in MCMC simulations to draw samples from $P(\theta \mid M)$.

MCMC simulations are performed using the Metropolis–Hastings algorithm with Gaussian jumping distribution. We use 8 parallel chains, each with 1.5 million steps from which the first 0.25 million are removed (burn-in). This gives altogether 10 million parameter samples from which we retain every thousandth sample, i.e. 10,000 samples for model simulations. We compared the between- and within-chain variance estimates using the r-hat convergence statistics to find r-hat < 1.01, suggesting a sufficient length of chains. The effective sample size of the final sample was found to be at least 1000.

For the biochemical cascade models, we assume wide log-normal priors and thus use log-parameters during MCMC sampling. After this reparameterization, we use a spherically symmetric Gaussian jumping distribution with independent variables and $\sigma = 0.2$ (in the log scale). For the mechanical model, we use an elliptical Gaussian jumping distribution with $\sigma = 1.0$ for masses, $\sigma = 0.2$ for damping coefficients, and $\sigma = 0.4$ for stiffness coefficients. For the final fit (trajectories of all three masses observed) we decrease step size by a factor of 4.

To minimise biases we select wide, non-informative priors, however, it is worth noting that in practical applications informative priors can significantly increase model performance.

In summary, our statistical model can be itemised as follows:

1. $\theta \sim$ Prior distribution.
2. $x(t) =$ Model ODEs solution for parameters θ and initial condition $x(0) = x_0$.
3. $y(t) = \mathsf{softlog}(x(t)) := \log(0.001 + x(t))$ (*only for signalling cascade*).
4. $m_{k,t,r} \sim \mathcal{N}(y_k(t), \sigma_{error})$.

where θ are the model parameters, σ_{error} is the measurement error and $m_{k,t,r}$ is the r-th measurement of variable k at time t, drawn from normal distribution $\mathcal{N}(y_k(t), \sigma_{error})$. Initial conditions for the ODE are kept constant and excluded from estimation. The *softlog* normalisation step is omitted for spring models. For all our models $\sigma_{error} = 0.3$.

Principal component analysis.   In Fig. 2F, Fig. S1 and Fig. S2D we use principal component analysis (PCA) to quantitatively describe the dimensionality reduction of the plausible parameter set. This approach works if the analysed set of points is "sufficiently flat". Points on a half circle would have two big PC values, even though the space is clearly one-dimensional. In such a case dimensionality can be investigated locally, as proposed by Little et al.[42]. In the half-circle example, this local PCA analysis would result in one small and one big PC value (for a sufficiently small spatial scale). We have verified that in our case local estimates result in higher PC values than the global ones; we thus conclude that our point clouds are sufficiently flat to justify using global PCA analysis.

Implementation.   We implemented our models in Python using jax[43] and diffrax[44] (for ODE integration), and the MCMC algorithm using blackjax[45].

Data availability

No experimental datasets were generated or analysed during the current study. All numerical results were obtained using code available at https://github.com/grfrederic/identifiability.

Code availability

Our source code is available at https://github.com/grfrederic/identifiability.

References

1. Brown, K. S. & Sethna, J. P. Statistical mechanical approaches to models with many poorly known parameters. *Phys. Rev. E* **68**, 021904 (2003).
2. Gutenkunst, R. N. *et al.* Universally sloppy parameter sensitivities in systems biology models. *PLoS Comput. Biol.* **3**, e189 (2007).
3. Raue, A. *et al.* Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics* **25**, 1923–1929 (2009).
4. Wieland, F.-G., Hauber, A. L., Rosenblatt, M., Tönsing, C. & Timmer, J. On structural and practical identifiability. *Curr. Opin. Syst. Biol.* **25**, 60–69 (2021).
5. Jaruszewicz-Błońska, J., Kosiuk, I., Prus, W. & Lipniacki, T. A plausible identifiable model of the canonical NF-κB signaling pathway. *PLoS ONE* **18**, e0286416 (2023).
6. Roosa, K. & Chowell, G. Assessing parameter identifiability in compartmental dynamic models using a computational approach: Application to infectious disease transmission models. *Theor. Biol. Med. Model.* **16**, 1 (2019).
7. Chowell, G. Fitting dynamic models to epidemic outbreaks with quantified uncertainty: A primer for parameter uncertainty, identifiability, and forecasts. *Infect. Dis. Model.* **2**, 379–398 (2017).
8. Transtrum, M. K. & Qiu, P. Model reduction by manifold boundaries. *Phys. Rev. Lett.* **113**, 098701 (2014).
9. Monsalve-Bravo, G. M. *et al.* Analysis of sloppiness in model simulations: Unveiling parameter uncertainty when mathematical models are fitted to data. *Sci. Adv.* **8**, 5952 (2022).
10. Brown, K. S. *et al.* The statistical mechanics of complex signaling networks: Nerve growth factor signaling. *Phys. Biol.* **1**, 184 (2004).
11. Cedersund, G. Conclusions via unique predictions obtained despite unidentifiability—New definitions and a general method. *FEBS J.* **279**, 3513–3527 (2012).
12. van Mourik, S., ter Braak, C., Stigter, H. & Molenaar, J. Prediction uncertainty assessment of a systems biology model requires a sample of the full probability distribution of its parameters. *PeerJ* **2**, e433 (2014).
13. Villaverde, A. F. & Banga, J. R. Reverse engineering and identification in systems biology: Strategies, perspectives and challenges. *J. R. Soc. Interface* **11**, 20130505 (2014).
14. Casey, F. P. *et al.* Optimal experimental design in an epidermal growth factor receptor signalling and down-regulation model. *IET Syst. Biol.* **1**, 190–202 (2007).
15. Liepe, J., Filippi, S., Komorowski, M. & Stumpf, M. P. H. Maximizing the information content of experiments in systems biology. *PLoS Comput. Biol.* **9**, e1002888 (2013).
16. Fox, Z. R. & Munsky, B. The finite state projection based Fisher information matrix approach to estimate information and optimize single-cell experiments. *PLoS Comput. Biol.* **15**, e1006365 (2019).
17. Lake, D., Corrêa, S. A. L. & Müller, J. Negative feedback regulation of the ERK1/2 MAPK pathway. *Cell Mol. Life Sci.* **73**, 4397–4413 (2016).
18. Kochańczyk, M. *et al.* Relaxation oscillations and hierarchy of feedbacks in MAPK signaling. *Sci. Rep.* **7**, 38244 (2017).
19. Hoffmann, A., Levchenko, A., Scott, M. L. & Baltimore, D. The IκB-NF-κB signaling module: Temporal control and selective gene activation. *Science* **298**, 1241–1245 (2002).
20. Lipniacki, T., Paszek, P., Brasier, A. R., Luxon, B. & Kimmel, M. Mathematical model of NF-κB regulatory module. *J. Theor. Biol.* **228**, 195–215 (2004).
21. Harris, S. L. & Levine, A. J. The p53 pathway: Positive and negative feedback loops. *Oncogene* **24**, 2899–2908 (2005).
22. Hat, B., Kochańczyk, M., Bogdał, M. N. & Lipniacki, T. Feedbacks, bifurcations, and cell fate decision-making in the p53 system. *PLoS Comput. Biol.* **12**, e1004787 (2016).
23. Toettcher, J. E., Weiner, O. D. & Lim, W. A. Using optogenetics to interrogate the dynamic control of signal transmission by the Ras/Erk module. *Cell* **155**, 1422–1434 (2013).
24. Nałęcz-Jawecki, P. *et al.* The MAPK/ERK channel capacity exceeds 6 bit/hour. *PLoS Comput. Biol.* **19**, e1011155 (2023).
25. Jeong, J. E. & Qiu, P. Quantifying the relative importance of experimental data points in parameter estimation. *BMC Syst. Biol.* **12**, 103 (2018).
26. Valderrama-Bahamóndez, G. I. & Fröhlich, H. MCMC techniques for parameter estimation of ODE based models in systems biology. *Front. Appl. Math. Stat.* **5**, 55 (2019).
27. Linden, N. J., Kramer, B. & Rangamani, P. Bayesian parameter estimation for dynamical models in systems biology. *PLoS Comput. Biol.* **18**, e1010651 (2022).
28. Hamon, J., Jennings, P. & Bois, F. Y. Systems biology modeling of omics data: Effect of cyclosporine a on the Nrf2 pathway in human renal cells. *BMC Syst. Biol.* **8**, 76 (2014).
29. Tavare, S., Balding, D. J., Griffiths, R. C. & Donnelly, P. Inferring coalescence times from DNA sequence data. *Genetics* **145**, 505 (1997).
30. Mikelson, J. & Khammash, M. Likelihood-free nested sampling for parameter inference of biochemical reaction networks. *PLoS Comput. Biol.* **16**, e1008264 (2020).
31. Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087–1092 (1953).
32. Geva-Zatorsky, N. *et al.* Oscillations and variability in the p53 system. *Mol. Syst. Biol.* **2**, 20060033 (2006).
33. Krishna, S., Jensen, M. H. & Sneppen, K. Minimal model of spiky oscillations in NF-kappaB signaling. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 10840–10845 (2006).
34. Machta, B. B., Chachra, R., Transtrum, M. K. & Sethna, J. P. Parameter space compression underlies emergent theories and predictive models. *Science* **342**, 604–607 (2013).
35. Raue, A., Kreutz, C., Theis, F. J. & Timmer, J. Joining forces of Bayesian and frequentist methodology: A study for inference in the presence of non-identifiability. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **371**, 20110544 (2013).
36. Fielding, M., Nott, D. J. & Liong, S.-Y. Efficient MCMC schemes for computationally expensive posterior distributions. *Technometrics* **53**, 16–28 (2011).
37. Conrad, P. R., Marzouk, Y. M., Pillai, N. S. & Smith, A. Accelerating asymptotically exact MCMC for computationally intensive models via local approximations. *J. Am. Stat. Assoc.* **111**, 1591–1607 (2016).

38. Cedersund, G. & Roll, J. Systems biology: Model based evaluation and comparison of potential explanations for given biological data. *FEBS J.* **276**, 903–922 (2009).
39. Daniels, B. C., Chen, Y.-J., Sethna, J. P., Gutenkunst, R. N. & Myers, C. R. Sloppiness, robustness, and evolvability in systems biology. *Curr. Opin. Biotechnol.* **19**, 389–395 (2008).
40. Sturm, O. E. *et al.* The mammalian MAPK/ERK pathway exhibits properties of a negative feedback amplifier. *Sci. Signal* **3**, 90 (2010).
41. Sigal, A. *et al.* Variability and memory of protein levels in human cells. *Nature* **444**, 643–646 (2006).
42. Little, A., Maggioni, M. & Rosasco, L. *Multiscale Geometric Methods for Estimating Intrinsic Dimension* (2011).
43. Bradbury, J. *et al. JAX: Composable Transformations of Python+NumPy Programs* (2018).
44. Kidger, P. *On Neural Differential Equations* (University of Oxford, 2021).
45. Lao, J. & Louf, R. *Blackjax: A Sampling Library for JAX* (2020).

Acknowledgements

Author contributions

F.G., P.N.J. and T.L. designed the study, F.G. wrote numerical codes, performed simulations and prepared figures. F.G. and T.L. with P.N.J.'s assistance wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-023-37939-8.

Correspondence and requests for materials should be addressed to T.L.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.