

Likelihood based observability analysis and confidence intervals for predictions of dynamic models

Clemens Kreutz,* Andreas Raue, and Jens Timmer
Freiburg Centre for Biosystems Analysis (ZBSA),
Freiburg Institute for Advanced Studies (FRIAS),
BIOSS Centre for Biological Signalling Studies,
Physics Department,
University of Freiburg
79104 Freiburg, Germany

Mechanistic dynamic models of biochemical networks such as Ordinary Differential Equations (ODEs) contain unknown parameters like the reaction rate constants and the initial concentrations of the compounds. The large number of parameters as well as their nonlinear impact on the model responses hamper the determination of confidence regions for parameter estimates. At the same time, classical approaches translating the uncertainty of the parameters into confidence intervals for model predictions are hardly feasible.

In this article it is shown that a so-called *prediction profile likelihood* yields reliable confidence intervals for model predictions, despite arbitrarily complex and high-dimensional shapes of the confidence regions for the estimated parameters. Prediction confidence intervals of the dynamic states allow a data-based observability analysis. The approach renders the issue of sampling a high-dimensional parameter space into evaluating one-dimensional prediction spaces. The method is also applicable if there are non-identifiable parameters yielding to some insufficiently specified model predictions that can be interpreted as non-observability. Moreover, a *validation profile likelihood* is introduced that should be applied when noisy validation experiments are to be interpreted.

The properties and applicability of the prediction and validation profile likelihood approaches are demonstrated by two examples, a small and instructive ODE model describing two consecutive reactions, and a realistic ODE model for the MAP kinase signal transduction pathway. The presented general approach constitutes a concept for observability analysis and for generating reliable confidence intervals of model predictions, not only, but especially suitable for mathematical models of biological systems.

I. INTRODUCTION

The major steps of the process of mathematical modeling comprise model discrimination, i.e. identification of an appropriate model structure, model calibration, i.e. estimation of unknown model parameters, as well as prediction and model validation. For all these topics it is essential to have appropriate methods assessing the certainty or ambiguity of any result for given experimental information.

For parameter estimation, there are several approaches to derive confidence intervals, like standard intervals which are based on an estimate of the covariance matrix of the parameter estimates [Sachs1984], bootstrap based confidence intervals [Davison1997, DiCiccio1987, Joshi2006], as well as likelihood based confidence intervals [Venzon1988, Raue2009]. For model discrimination, significance statements can be obtained by statistical tests. However, for model predictions, there are still demands for methodology that is applicable for mathematical models like ODEs used to describe the dynamics of a system in a variety of scientific fields e.g. in molecular biology [Hlavacek2009, Swameye2003], but also in medical

research, chemistry, engineering, and physics.

The mere estimation of parameters is often not the final aim of an investigation. More frequently, it is desired to utilize the parametrized model to generate model predictions such as the dynamic behavior of unobserved components. Classically, the uncertainty in the model parameters is attempted to be translated into corresponding prediction confidence intervals. For models that depend linearly on the model parameters, as it occurs in classical regression models, this is well studied and known as propagation of uncertainty based on standard errors. This approach is appropriate and sufficient for many applications. However, e.g. for biochemical networks, the model responses depend nonlinearly on the model parameters. Here, the boundaries of the parameter confidence region can exhibit arbitrarily complex shape and are usually difficult to translate into boundaries for the prediction confidence intervals. Therefore, established approaches aim to scan the entire parameter subspace which is in a sufficient agreement with the experimental data to propagate the parameters confidence regions into confidence intervals for the model predictions. The major challenge is the complex nonlinear interrelation between parameters and model responses which requires that the parameter space has to be densely sampled to capture all scenarios of model predictions. For models with tens to hundreds of parameters this is numerically demanding or even infeasible because high

*Electronic address: ckreutz@fdm.uni-freiburg.de

dimensional spaces cannot be densely sampled. This is an issue often referred to as the curse of dimensionality in literature [Marimont1979, Scott1991].

There are several methods for an approximate sampling of the parameter space, e.g. the Markov Chain Monte Carlo (MCMC) methods [Gelman2003, Kass1998], and bootstrap based approaches [Joshi2006, Molinaro2005]. However, for the ODE models used to describe interaction networks, these methods are numerically very demanding and provide only very rough approximations. Therefore, it is difficult to control the coverage of the prediction confidence intervals for these approaches. Moreover non-identifiable parameters are not explicitly considered hampering the convergence of these sampling techniques and yielding results that are questionable and difficult to interpret [Bayarri2004].

Conceptually related to the prediction profile likelihood approach presented here, [Raue2009, Raue2010] presented an approach for the determination of confidence intervals for the model parameters by sampling the parameter profile likelihood. For their approach the problems concerning identifiability are resolved and the computation is feasible also for high-dimensional parameter spaces. Nevertheless, the direct translation to confidence intervals of the model trajectories only works as an approximation yielding a coverage rate that is sometimes lower than desired.

The idea of the prediction profile likelihood presented here is to determine prediction confidence without an explicit sampling strategy for the parameter space. Instead, a certain fixed value for a prediction is used as a nonlinear constraint and the parameter values are chosen via constraint optimization of the likelihood. This does neither require a unique solution in terms of identifiability nor confidence intervals for the parameter estimates. The constraint maximum likelihood approach checks the agreement of a predicted value with the experimental data. By repeating this procedure for continuous variations of the predicted value, a *prediction profile-likelihood* is obtained. Thresholding the prediction profile likelihood yields statistically accurate confidence intervals. The desired level of confidence which coincides with the level of agreement with the experiments is controlled by the threshold.

The theoretical background of the prediction profile likelihood, also called *predictive likelihood* has been already studied [Hinkley1979]. Moreover, related ideas are already applied in the context of generalized linear mixed models [Booth1998], unobserved data points [Butler1986]. The linear approximation has been applied in nonlinear regression analyses [Cooley1989]. A review of prediction profile likelihood approaches and a modification to sufficiency-based predictive likelihood is provided in [Bjornstad1990].

In this paper, this concept is applied to ODE models occurring in mechanistic dynamic models, e.g. in molecular and systems biology. In this context the approach allows a data-based observability analysis. Moreover, it

is extended to obtain confidence intervals for validation experiments.

II. RESULTS

A. Small illustration model

First, a small but illustrative model of two consecutive reactions



with rates $\theta_1 = 0.05, \theta_2 = 0.1$ and initial conditions $A(0) = \theta_3 = 1, B(0) = 0, C(0) = 0$ is utilized to illustrate our approach. For this purpose, it is assumed that $C(t)$ is measured at $t = 0, 10, \dots, 100$.

For the simulated measurements, Gaussian noise $\varepsilon \sim N(0, \sigma^2)$ with $\sigma = 0.1$ has been assumed which corresponds to a typical signal-to-noise ratio for applications in cell biology of around 10%. If an experimental setup would not allow for negative measurements, a log-normal distribution of the observational noise could be more appropriate. Then, the Gaussian setting is obtained after a log-transformation of the data [Kreutz2007]. Such transformations and preprocessing procedures would have to be performed before the analysis starts.

Panel (a) in Fig. 1 shows the dynamics of $A(t)$, $B(t)$, and $C(t)$ as well as a typical noise realization. Such simulated data realizations are utilized to calculate the prediction- and validation profile likelihood for the dynamic states.

Panel (b) shows, as an example, the prediction profile likelihood and the validation profile likelihood for the same noise realization for predicting $A(t)$ at time point $t = 10$. The validation profile likelihood has been calculated for validation data with 10% measurement noise, as it was assumed for the measurements. The vertical axis is minus two times the log-likelihood which corresponds to the residual sum of squares. For illustration purposes, the minimum of the log-likelihood LL^* is shifted to zero in all figures. Three thresholds corresponding to 68%, 90%, and 99% confidence levels are plotted as horizontal lines. As explained in the section ‘Materials and Methods’, the projections to the horizontal axis yields the respective confidence intervals for a prediction or for a validation experiment. The constraint optimization procedure is infeasible for $A(t) \leq 0$ and therefore the PCIs automatically account for strictly positive values of A .

The calculation of the prediction and validation confidence intervals has been repeated for $t = 0, 10, \dots, 100$ and all three dynamic states $A(t)$, $B(t)$, $C(t)$. In panels (c)-(e), the respective prediction confidence intervals (PCIs) are plotted as well as the prediction profile likelihood. The corresponding validation profile likelihood functions and the respective validation confidence intervals are shown in the supplemental information (SI). For plotting the confidence intervals along the time axis, the

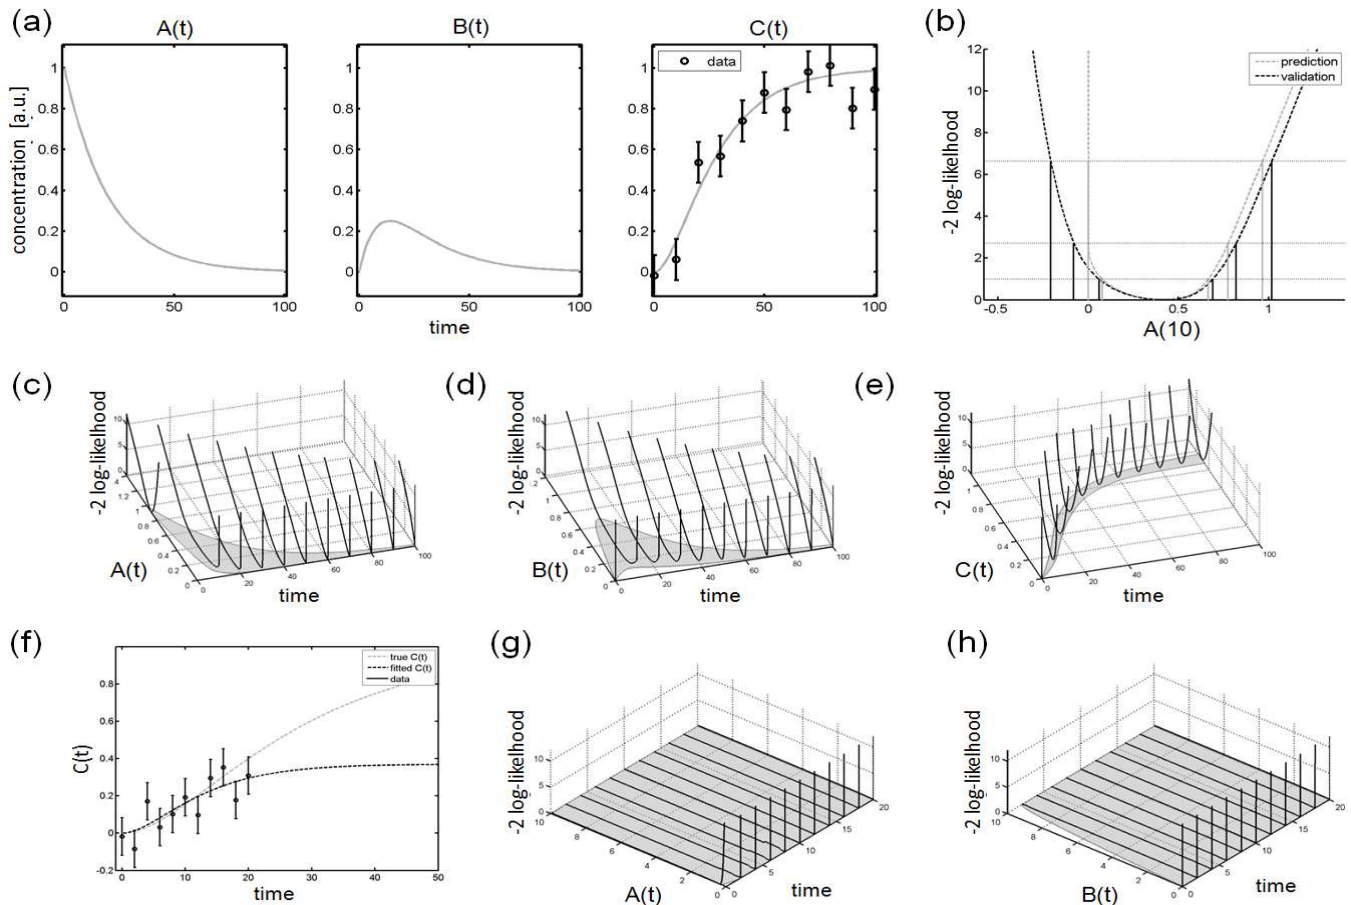


FIG. 1: The three figures in panel (a) show the dynamics and measurement realization for the small model used for illustration purpose. $C(t)$ is measured and the dynamics of all states, i.e. $A(t)$, $B(t)$, and $C(t)$, is intended to be predicted. Panel (b) shows as an example the prediction profile likelihood (gray dashed curve) and validation profile likelihood (black dashed curve) of $A(t=10)$. Thresholding yields confidence intervals for prediction (gray vertical lines) and validation (black vertical lines). The three thresholds and the respective projections correspond to $\alpha=68\%$, 90% , and 99% confidence intervals. The VCIs are larger than the PCIs, because they account for the measurement error of a validation data point. Panels (c)-(e) show prediction confidence intervals (gray) for the unobserved states $A(t)$, $B(t)$, as well as for the measured state $C(t)$. The prediction profile likelihood functions are plotted as black curves in vertical direction. Non-observability is illustrated in panels (f)-(h). Panel (f) shows a realization of the measurements for a design which does not provide sufficient information about the steady state of C . This leads to a flat prediction profile likelihood for large values for $A(t)$ as shown in panel (g), as well as for $B(t)$ for $t>0$ as plotted in panel (h). A flat prediction profile likelihood in turn yields unbounded prediction and validation confidence intervals and non-observability of $A(t)$ and $B(t)$ as indicated by the gray shaded regions.

PCIs evaluated the eleven time points have been interconnected by cubic piecewise interpolation. The displayed confidence intervals constitute the propagation of information from the measurements of $C(t)$ to predictions of the dynamics of the compound concentrations. Because C is the measured compound in our example, the prediction confidence intervals for C are much smaller than for A and B . However, also A and B yield bounded prediction confidence intervals which can be interpreted as *observability* of these dynamic states.

In the Appendix, the reliability of our confidence intervals is investigated by calculating the coverage, i.e. by comparing the confidence level with the frequency that the true value is inside the prediction confidence inter-

val. For our example, it will be demonstrated that confidence intervals using the asymptotic threshold sometimes yield slightly conservative intervals. Also an algorithm to improve the threshold is provided which yields slightly smaller confidence intervals with the correct coverage.

To illustrate the effect of *non-observability*, the assumption about the available experimental information is slightly changed. The measurements are simulated for earlier and closer time steps, i.e. for $t = 0, 2, \dots, 20$. Panel (f) of Fig. 1 shows that these time points sample only the transient increase of $C(t)$. Hence, such a design does not provide sufficient information about the steady state level of C . In other words, the modification limits the available information about the total amounts

of the compounds. This, in turn, makes $A(t)$ and $B(t)$ non-observable.

Panel (g) shows the prediction confidence intervals for $A(t)$. In the chosen setting, the predictions are unbounded towards infinity and therefore $A(t)$ is non-observable. In panel (h), it is also shown that $B(t)$ is non-observable. According to the model definition, $B(0)$ is known to be zero, but for $t > 0$, unbounded prediction confidence intervals are obtained which indicate non-observability of $B(t)$.

B. MAP kinase signaling model

Next, an ODE model of cellular signal transduction has been used to illustrate our method in a realistic setting. For this purpose, a model of the *mitogen-activated protein (MAP) kinases* which is one of the most extensively studied signal transduction pathway, is utilized. The chosen model [Kholodenko2000] consists of eight dynamic states describing the time dependency of the MAP kinases Raf, Raf*, Mek, Mek*, Mek**, Erk, Erk*, and Erk** which play a very prominent role in many cellular processes, e.g. in cell proliferation. A star '*' denotes phosphorylation of the protein which biologically acts as activation.

Panel (a) in Fig. 2 provides a summary of the MAP kinase signaling pathway. The enzymatic reactions in the ODE model are described as Michaelis-Menten rate equations, i.e. each reaction is parametrized with a maximal enzymatic rate and a Michaelis constant. As in the original publication, the parameters of the two consecutive phosphorylation and dephosphorylation steps of Mek and Erk are assumed to be identical and the initial concentrations are assumed to be known. In this setting, 14 parameters are estimated out of three times eleven data points. Details about the model are provided in the Appendix.

It is assumed that the total amount of the phosphorylated forms for each protein, i.e. Raf*, the sum of Mek* and Mek** as well as the sum of Erk* and Erk**, are measured. This observational assumption holds for example for phospho-specific antibodies such as utilized for Western blotting. The measurement times are set to 0, 100, ..., 1000 seconds. Again, additive Gaussian noise is assumed. The standard deviation has been set to $\sigma = 10$ nM.

In panel (b) of Fig. 2 a typical noise realization is displayed. Panel (c) shows the prediction confidence intervals (dark gray) and the validation confidence intervals (light gray) for this noise realization calculated for all dynamic states. The size of the confidence intervals is plotted as a dashed-dotted line.

The prediction confidence intervals show how precisely the dynamics is inferred by the available data. The temporal behavior of Raf, Raf* is quite well determined, i.e. the size of the PCI is below 40 nM. Similarly, the unphosphorylated states of Mek and Erk have narrow prediction confidence intervals. For Mek* the concen-

tration dynamics is only predicted within rather large intervals which for most time points nearly span a range between zero and 100 nM.

The largest absolute size of the prediction confidence interval of 176 nM is obtained for Erk** after 181 seconds. This is the point in time where the negative feedback is activated. Additional experimental investigation of this condition is very informative to further specify the dynamic behavior of the MAP kinase cascade in our example. Further considerations concerning experimental planning are provided in detail in the Appendix.

III. DISCUSSION

Existing approaches for prediction confidence intervals like MCMC [Marjoram2003] or bootstrap procedures are based on forward evaluations of the model for many parameter values. This works reasonably well for a low dimensional parameter space and if the target density function, i.e. the parameter space to be sampled, is well-behaved [Bayarri2004]. However, sampling nonlinear high-dimensional spaces densely is impractical and it is almost impossible to ensure that sampling the parameter space covers all prediction scenarios. Especially in biological applications the target distributions frequently inherit strong and nonlinear functional relations. In the case of non-identifiability, the parameter space to be sampled is not restricted rendering convergence near to impossible.

In this paper, we applied a contrary procedure. The model prediction space is sampled directly and the corresponding model parameters are determined by constraint maximum likelihood to check the agreement of the predictions with the data. This concept yields the prediction profile likelihood which constitutes the propagation of uncertainty from experimental data to predictions.

If a comprehensive prior, i.e. for all parameters, would be available, a Bayesian procedure like MCMC where marginalization, i.e. integration over the nuisance dimensions is feasible could have superior performance. However, in cell biology applications, prior knowledge is very restricted because kinetic rates and concentrations are highly dependent on the cell type and biological context, e.g. on the cellular environment and biochemical state of a cell. Therefore, there is usually at most some prior information for few parameters available. Such prior information can be incorporated in our procedure without restricting its applicability by generalizing maximum likelihood estimation to maximum a-posterior estimation as discussed in the Appendix.

In general, generating prediction confidence intervals given the uncertainty in the high-dimensional nonlinear parameter space requires large numerical efforts. However, this complication primarily originates from the complexity of the issue itself rather than from the methodological choice. In fact, the aim is approached by the prediction profile likelihood in a very efficient

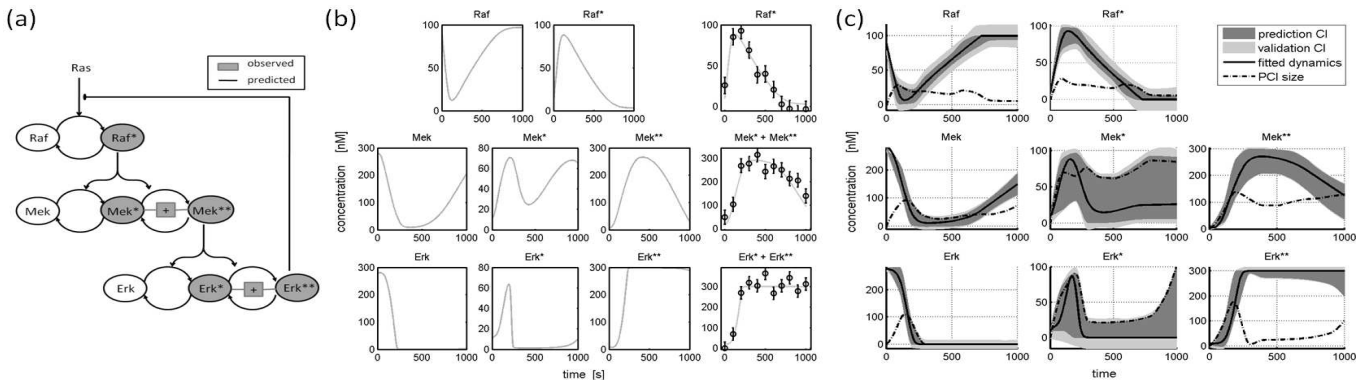


FIG. 2: Panel (a) shows the MAP kinase model according to [Kholodenko2000]. It is assumed that the phosphorylated compounds are measured. The dynamics of all compounds is intended to be predicted to illustrate the prediction profile likelihood approach. In panel (b) the dynamics of the MAP kinase model as well as simulated data set are plotted. The 90% confidence intervals of the dynamic variables for predictions (dark gray) and for validation experiments (light gray) for this noise realization are plotted in panel (c). The size of the PCI is plotted as a dashed-dotted line. In absolute concentrations, the dynamics of Erk** has the largest PCI at $t=181$ seconds, i.e. when the negative feedback is activated. Also, the dynamics of Mek* is only badly observable in our example. Measurements of both would be very informative for better calibrating the model.

manner because scanning the parameter space by the constrained optimization procedure to explore the data-consistent predictions is more efficient than sampling parameter space without considering the predictions like it is performed for MCMC. Instead of sampling a high-dimensional parameters space, only the prediction space has to be explored for calculating a prediction profile likelihood, i.e. the optimization of the parameters reduces the high-dimensional sampling problem to exploring a single dimension. For the parameter profile likelihood, it has been demonstrated [Raue2009] that the computational effort only scales slightly super-linear with the number of parameters. This result does, due to the similarity of the computations, carry over to the prediction profile likelihood.

The prediction confidence regions introduced above has to be interpreted *point-wise*. This means that a confidence level α controls errors of type 1 which is the probability that the model response for the true parameters is inside the prediction confidence interval for a single prediction condition if many realizations of the experimental data and the corresponding prediction confidence intervals are considered.

In contrast, if a single data set is utilized to generate many prediction intervals, e.g. predictions for several points in time as performed above, the results are *statistically dependent*, i.e. the realization of the PCI of a neighboring time point is very similar and therefore correlated. Therefore, the prediction confidence intervals for a compound for two adjacent points in time very likely both contain the true value, or neither. In such an example, a *common* prediction confidence region for two statistically dependent predictions would require a two-dimensional prediction profile likelihood. This topic, however, is beyond the scope of this article.

The prediction profile likelihood also provides a concept for experimental planning. Experimental conditions with a very narrow prediction confidence interval are very accurately specified by the available data. New measurements for such a condition on the one hand does not provide very much additional information to better calibrate the model parameters, and hence is from this point of view a bad choice for additional measurements. On the other hand, it very precisely predicts the model behavior under these certain conditions and is therefore a very powerful candidate setting for validating the model structure. Contrarily, large prediction confidence intervals indicate conditions which are weakly specified by the existing data and therefore constitute informative experimental designs for better calibrating the model. Because a design optimization on the basis of the prediction profile likelihood does not require any linearity approximation like common experimental design techniques, e.g. based on the *Fisher information* [Kreutz2009], the presented procedure could be very valuable for ODE models which are typically highly nonlinear.

Another potential of the prediction profile likelihood shown in this article is its interpretation in terms of *observability*. This term is very commonly used in control theory to characterize whether the dynamics of some unobserved variables can be inferred by the set of feasible experiments. The theory in this field is based on analytical calculations, i.e. the limited amount and inaccuracy of the data is usually not considered. In this article, it has been shown that the prediction profile likelihood allows for a general data-based approach to check whether there is enough information about unobserved dynamic states in the given experimental design and realization of measurements. Therefore, in analogy to the terminology of *practical identifiability* [Raue2009], we would suggest

to term observability for a given data set, i.e. a restricted prediction confidence interval, as *practical observability*.

Finally, it should be noted, that a prediction could be any function of the compounds and the parameters. In applications, e.g. a ratio of two compound concentrations is a characteristics of interest. In principle also integrals, peak positions and other functions of the dynamic states can be considered as predictions which could be targets for observability considerations as well as for the calculation of prediction and validation confidence intervals.

IV. SUMMARY

Generating model predictions is a major task in mathematical modeling. For the dynamic mechanistic models as they are applied e.g. in molecular and systems biology, the confidence regions from parameter estimation can have arbitrarily complex shapes. Therefore, it is very difficult or even impossible to sample the parameter space appropriately to generate confidence intervals for predictions. This in turn impedes a data-based observability analysis for the dynamic states.

In this article, the prediction profile likelihood approach is presented as a methodology which directly calculates the set of model predictions which are consistent with existing measurements. This concept constitutes a powerful tool for assessing model predictions, performing observability analyses, and experimental design. The method is feasible for arbitrary dimensions of the parameter space. It only requires a proper calculation of the maximum likelihood value, i.e. a numerically working parameter optimization procedure. The task of sampling a high-dimensional parameter space reduces to scanning a one-dimensional prediction space. It therefore allows the calculation of confidence intervals for model predictions as well as confidence intervals for the outcome of validation experiments.

The applicability of the approach has been shown by a small but instructive system of two consecutive reactions and a published model for MAP kinase signaling. For the small system, it has been shown that the prediction profile likelihood yields desired coverage properties. Moreover, a setting inducing non-observability has been investigated which is characterized by unbounded prediction confidence intervals. For the MAP kinase model, prediction confidence intervals and validation confidence intervals for all dynamic states have been determined on the basis of measurements of the phosphorylated proteins. In addition, the applicability of the approach for experimental planning has been demonstrated.

V. METHODOLOGY

A. The prediction profile likelihood

For additive Gaussian noise $\varepsilon \sim N(0, \sigma^2)$ with known variance σ^2 , two times the negative log-likelihood

$$-2 \text{LL}(y|\theta) = \sum_i \frac{(y_i - F(t_i, u, \theta))^2}{\sigma^2} + \text{const.} \quad (2)$$

of the data y for the parameters θ is except a constant offset identical to the residual sum of squares $\text{RSS}(\theta) = \sum_i (y_i - F(t_i, u, \theta))^2 / \sigma^2$. In this case, maximum likelihood estimation is equivalent to standard least-squares estimation

$$\hat{\theta} = \arg \max_{\theta} \text{LL}(y|\theta) \equiv \arg \min_{\theta} \text{RSS}(\theta), \quad (3)$$

i.e. to minimizing the residual sum of squares. $F = g(x(t, u, \theta), \theta)$ denotes the model response which is in our case given after integration of a system of differential equations

$$\dot{x}(t) = f(x(t), u(t), \theta) \quad (4)$$

with an externally controlled input function u and a mapping to experimentally observable quantities

$$y(t) = g(x(t), \theta) + \varepsilon(t). \quad (5)$$

The parameter vector θ comprises the kinetic parameters as well as the initial values, and additional offset or scaling parameters for the observations.

It has been shown [Raue2009] that the profile likelihood

$$\text{PL}(\theta_i) = \max_{\theta_{j \neq i}} \text{LL}(\theta|y) \quad (6)$$

for a parameter θ_i given a data set y yields reliable confidence intervals

$$\text{CI}_{\alpha}(\theta_i|y) = \{\theta_i \mid -2\text{PL}(\theta_i) \leq -2\text{LL}(y)^* + \text{icdf}(\chi_1^2, \alpha)\} \quad (7)$$

for the estimation of a single parameter. Here, α is the confidence level and $\text{icdf}(\chi_1^2, \alpha)$ denotes the α quantile of the chi-square distribution with one degree of freedom which is given by the respective inverse cumulative density function. LL^* is the maximum of the log-likelihood function after all parameters are optimized. In [6], the optimization is performed for all parameters except θ_i . The analogy of likelihood-based parameter and prediction confidence intervals is discussed in the Appendix. The desired coverage

$$\text{Prob}(\theta_i \in \text{CI}_{\alpha}(\theta_i)) = \alpha, \quad (8)$$

i.e. the probability that the true parameter value is inside the confidence interval, holds for [7] if the magnitude of the residual sum of squares by fitting of θ_i is χ_1^2 distributed. This is given asymptotically as

well as for linear parameters and is a good approximation under weak assumptions [Feder1968, Seber1989]. If the assumptions are violated, the distribution of the magnitude of the decrease has to be generated empirically, i.e. by Monte-Carlo simulations, as discussed in the Appendix.

The *experimental design* $D = \{t, g, u\}$ comprises all environmental conditions which can be controlled by the experimenter like the measurement times t , the observables g , and the input functions u . A *prediction* $z = F(D_{\text{pred}}, \theta)$ is the response of the model F for a prediction condition $D_{\text{pred}} = \{t_{\text{pred}}, g_{\text{pred}}, u_{\text{pred}}\}$ specifying a prediction observable g_{pred} evaluated at time point t_{pred} given the externally controlled stimulation u_{pred} .

In some cases the observable g_{pred} corresponds to measuring a dynamic variable $x(t)$ directly, i.e. it corresponds to a compound whose concentration dynamics is modeled by the ODEs. In a more general setting the observable is defined by an observational function $g_{\text{pred}}(x(t), \theta)$ depending on several dynamic variables x . Therefore, g_{pred} does neither have to coincide with a dynamic variable nor with an observational function g of the measurements performed to build the model.

In analogy to [8], the desired property of a prediction confidence interval $\text{PCI}_\alpha(D|y)$ derived from an experimental data set y with a given significance level α is that the probability

$$\text{Prob}(F(D_{\text{pred}}, \theta_{\text{true}}) \in \text{PCI}_\alpha(D|y)) = \alpha \quad (9)$$

that the true value of $F(D_{\text{pred}}, \theta_{\text{true}})$ is inside the prediction confidence interval PCI_α is equal to α . In other words, the PCI covers the model response for the true parameters with a proportion α of the noise realizations which would yield different data sets y .

The prediction profile likelihood

$$\text{PPL}(z) = \max_{\theta \in \{\theta | F(D_{\text{pred}}, \theta) = z\}} \text{LL}(y|\theta) \quad (10)$$

is obtained by maximization over the model parameters satisfying the constraint that the model response $F(D, \theta^*)$ after fitting is equals to the considered value z for the prediction. The prediction confidence interval is in analogy to [7] given by

$$\text{PCI}_\alpha(D_{\text{pred}}|y) = \{z | -2\text{PPL}(z) \leq -2\text{LL}^*(y) + \text{icdf}(\chi_1^2, \alpha)\} \quad (11)$$

i.e. the set of predictions $z = F(D_{\text{pred}}, \theta)$ for which the PPL is below a threshold given by the χ_1^2 distribution. In analogy to likelihood based confidence intervals for parameters, such PCI yields the smallest unbiased confidence intervals for predictions for given coverage α [Cox1994].

Instead of sampling a high-dimensional parameter space, the prediction profile likelihood calculation comprises sampling of a one-dimensional prediction space by evaluating several predictions z . Evaluating the maximum of the likelihood satisfying the prediction constraint does in general not require an unambiguous point in the parameter space as in the case of structural non-identifiabilities.

In analogy to profile likelihood for parameter estimates, the significance level determines the threshold for the PPL, which is given asymptotically by the quantiles [7] of the χ_1^2 distribution [Meeker1995]. In the Appendix, a Monte-Carlo algorithm is presented which can be used to calculate the threshold in cases where the asymptotic assumption is violated.

B. The validation profile likelihood

Likelihood-based confidence interval like [7] or [11] correspond to the region where a likelihood ratio test would not reject the model. Having a prediction confidence interval, the question arises whether a model has to be rejected if a validation measurement is outside the predicted interval. This, in fact, would hold if a “perfect” validation measurement would be available, i.e. a data point without measurement noise. For validation experiments, however, the outcome is always noisy and is therefore expected to be more frequently outside the PCI than the true value. Hence, the prediction confidence interval [11] has to be generalized for application to a validation experiment.

For a validation experiment, we therefore introduce a *validation profile likelihood* VPL and a corresponding *validation confidence interval* $\text{VCI}_\alpha^{\text{SD}}$ in the following. In such a setting, a confidence interval should have a coverage

$$\text{Prob}(z \in \text{VCI}_\alpha^{\text{SD}}(D_{\text{vali}}|y)) = \alpha \quad (12)$$

for the validation data point $z \sim N(\mu, \text{SD}^2)$ with expectation $\mu = F(D_{\text{vali}}, \theta_{\text{true}})$ and variance SD^2 . Here, D_{vali} denotes the design for the validation experiment. A validation confidence interval satisfying [12] allows a rejection of the model if a noisy validation measurement with error SD is outside the interval.

$\text{VCI}_\alpha^{\text{SD}}$ for validation data can be calculated by relaxing the constraint [10] used to compute the prediction profile likelihood. Because in this case, the model prediction does not necessarily have to coincide with the data point z . Instead, the deviation from the validation data point is penalized equivalently to the data y . The agreement of the model with the data y and the validation measurement z is then given by

$$\text{LL}(z, y|\theta) = \sum_i \left(\frac{y_i - F(D_i, \theta)}{\sigma} \right)^2 + \left(\frac{z - F(D_{\text{vali}}, \theta)}{\text{SD}} \right)^2 \quad (13)$$

We now define the validation profile (log-)likelihood

$$\text{VPL}^{\text{SD}}(z|y) = \text{LL}^*(z, y) = \text{LL}(z, y|\theta^*) \quad (14)$$

with $\theta^* = \theta^*(z, y) = \arg \max_\theta \text{LL}(z, y|\theta)$ as the maximized joint log-likelihood in [13] read as a function of z . The corresponding validation confidence interval is given by

$$\text{VCI}_\alpha^{\text{SD}}(D_{\text{vali}}|y) = \{z | -2\text{VPL}^{\text{SD}}(z|y) \leq -2\text{LL}^*(z, y) + \text{icdf}(\chi_1^2, \alpha)\} \quad (15)$$

Optimization of the likelihood [13] minimizes both, the contribution of the data $\text{RSS}(y)$, and the mismatch with the fixed prediction value z . The model response $F(D_{\text{pred},\theta^*})$ obtained after this parameter optimization can be interpreted as a prediction z' satisfying the constraint optimization problem [10] considered for the prediction profile likelihood. It holds

$$\text{LL}^*(z, y; \text{SD} > 0) - \frac{1}{2} \frac{(z - F(D_{\text{vali}}, \theta^*))^2}{\text{SD}^2} = \text{LL}^*(z', y; \text{SD} = 0), \quad (16)$$

i.e. the validation profile likelihood LL^* can be scaled to the prediction likelihood via

$$\text{PPL}(z'|y) = \text{VPL}^{\text{SD}}(z|y) - \frac{1}{2} \frac{(z' - z)^2}{\text{SD}^2} \quad (17)$$

where $z' = F(\theta^*(z, y, \text{SD} > 0))$ is the model response for θ^* estimated from z and y .

Optimization with nonlinear constraints is a numerically challenging issue. Therefore, [17] provides a helpful way to omit constraint optimization. The VPL can be calculated with $\text{SD} > 0$ like a common least-squares minimization and is then afterwards rescaled to obtain the PCI for the true value.

Acknowledgments

The authors thank our long-term experimental collaboration partners, especially Dr. Maria Bartolome-Rodriguez and Prof. Ursula Klingmüller and their groups for their support and their experience in practically relevant issues. In addition, the authors acknowledge financial support provided by the bmbf-grants 0315766-VirtualLiver, 0315415E-LungSys and 0313921-FRISYS.

Appendix A: Supplementary information

1. Re-parametrization

Parameter estimation, i.e. the prediction of a parameter value out of experimental data, can be seen as a special case of a model prediction. Then, the parameter profile likelihood coincides with the prediction profile likelihood and the respective parameter confidence intervals correspond to prediction confidence intervals. In this sense, the prediction profile likelihood generalizes the parameter profile likelihood. In fact, the idea of the prediction profile likelihood and the calculation of prediction confidence intervals, e.g. the choice of the threshold, is very intuitive for this special case.

In other situations, an analog strategy would require a re-parametrization of the model in a way that the desired model prediction is unambiguously given by the value of a single new parameter. Then, again the profile likelihood for the new parameter would give a confidence interval for the prediction. In this case, without loss of generality such a parameter can be denoted by θ'_1 . Then, the re-parametrization would be a transformation

$$T : \{\theta_1, \dots, \theta_{n_p}\} \rightarrow \{\theta'_1, \dots, \theta'_{n_p}\} \quad (\text{A1})$$

of the n_p parameters θ to new parameters $\theta'_1, \dots, \theta'_{n_p}$ where all predictions for the condition D_{pred} satisfy

$$F'(D_{\text{pred}}, \theta') = F'(D_{\text{pred}}, \theta'_1) . \quad (\text{A2})$$

Here, $F' = F \circ T^{-1}$ denotes the model for the transformed parameters. For a transformation satisfying [A2], any change of the parameters $\theta'_2, \dots, \theta'_{n_p}$ would not affect F' , because T is chosen in a way that the effect of $\theta_2, \dots, \theta_{n_p}$ is orthogonal to the effects of θ_1 .

However, because ODE systems can only be solved analytically for special cases, such a re-parametrization cannot be found explicitly for most realistic models. This restriction can be resolved numerically by an implicit re-parametrization which is obtained by a constrained nonlinear optimization procedure. This idea yields the prediction profile likelihood

$$\text{PPL}(z) = \max_{\theta \in \{\theta | F(D_{\text{pred}}, \theta) = z\}} \text{LL}(y|\theta) \quad (\text{A3})$$

which is obtained by maximization over the model parameters satisfying the constraint that the model response $F(D, \theta^*)$ after fitting is equals to the considered value z for the prediction. In this case, the 'new parameter' is the predicted value itself, i.e. $z \equiv \theta'_1$ and F' is the identity function.

Equation [A3] also resolves the formal issue which occurs if there is not a unique parameter set θ given by the constraint $F(D_{\text{pred}}, \theta) = z$. If there are several such parameter sets, the ambiguities either vanish by taking the parameter set with maximize the log-likelihood, or they are not relevant because only the value of the maximized log-likelihood enters the calculation.

Fit the model M to the measurements obtained for the design D . This model is considered as the 'truth'.

Define the prediction design D_{pred} .
Evaluate true model prediction $z_{\text{true}} = F(D_{\text{pred}})$.

for a noise realization ε_i
 $y_i(D) = \text{SIMULATE_DATA}(M, D, \varepsilon_i)$
 $\text{PPL}_i^* = \text{CALIBRATE_MODEL}(y_i(D), M)$
 $\text{PPL}_i(z_{\text{true}}) = \text{CALIBRATE_MODEL_CONSTRAINT}(y_i(D), M, z_{\text{true}})$
 $\text{epdf}_i = \text{PPL}_i(z_{\text{true}}) - \text{PPL}_i^*$
 $\text{threshold}(\alpha) = \text{quantile}_{\varepsilon_\alpha}(\text{ecdf})$

FIG. 3: A Monte-Carlo algorithm for calculating the profile likelihood threshold empirically. New noise realizations $y_i(D)$ are utilized to calculate the distribution of $\text{PPL}_i(z_{\text{true}}) - \text{PPL}_i^*$. The α quantile of this distribution can be used as a threshold for prediction confidence intervals instead of the asymptotic threshold, i.e. instead of the α quantile of the χ^2_1 distribution.




FIG. 4: The Monte-Carlo approach allows a comparison of the asymptotic thresholds with the empirically calculated, i.e. the correct thresholds. Here, the thresholds corresponding to 0.05, 0.1, ..., 0.95, 0.99 confidence levels have been plotted for nine different prediction scenarios. In our example, the asymptotic thresholds are slightly too large for predictions of $A(t)$ and $B(t)$ which makes the asymptotic confidence intervals conservative.

2. Profile likelihood threshold

A suitable parameter transformation makes the prediction profile likelihood equivalent to the parameter profile likelihood. Therefore, the following discussion holds for both, for parameter and for prediction confidence intervals.

In general, fitting a model to experimental data reduces the residual sum of squares RSS. In the asymp-

otic case, i.e. for a large number of data points, it can be shown that the decrease of RSS due to fitting one parameter is chi-square distributed with one degree of freedom. This result also holds exactly in the non-asymptotic case for linear parameters. This outcome is utilized to define the asymptotic threshold for profile likelihood confidence intervals [Raue2009].

For nonlinear parameters, the distribution of the decrease of the residual sum of squares by the parameter estimation procedure, has not yet been derived for the general setting. However, since the profile likelihood based confidence intervals are independent on bijective transformations of the parameter space [Cox1994], the assumption also holds if there is such a transformation, which makes the parameter of interest linear at least within its confidence interval. Such a transformation only has to exist, it is not required to derive it analytically.

A situation where such a transformation does not exist occurs if the nonlinearity yields a non-monotone dependency of the profile likelihood, i.e. there are several local minima in the confidence interval. In this case, there is a larger decrease of the residual sum of squares and the standard threshold yields conservative results, i.e. the calculated confidence intervals are too large for the desired confidence level α .

In Fig. 3, a procedure is presented for checking the standard threshold. It is a Monte-Carlo analysis of the impact of the nonlinear constraint used to calculate the prediction profile likelihood on the magnitude of overfitting. In Fig. 4, the asymptotic thresholds corresponding to $\alpha = 0.05, 0.1, \dots, 0.95, 0.99$, i.e. the quantiles of the χ_1^2 distribution, are compared with the empirically calculated thresholds for several prediction scenarios. For the model

$$A \rightarrow B \rightarrow C \quad (\text{A4})$$

the asymptotic thresholds are slightly too large for predicting $A(t)$ and $B(t)$ on the basis of measurements of $C(t)$. This makes the asymptotic confidence intervals conservative. The impact on the coverage is discussed in the following section.

3. Coverage

The coverage

$$C = \text{Prob}(F(D_{\text{pred}}, \theta_{\text{true}}) \in \text{PCI}_\alpha(D|y)) \quad (\text{A5})$$

is the probability that the $\text{PCI}_\alpha(z|y)$ contains the true value $F(D_{\text{pred}}, \theta_{\text{true}})$. A desired property of any confidence interval is that the coverage coincides with the confidence level α .

Fig. 5 shows the estimated coverage of the prediction confidence intervals calculated for nine different prediction scenarios. In these scenarios $A(2)$, $B(2)$, $C(2)$, $A(10)$, $B(10)$, $C(10)$, $A(50)$, $B(50)$, $C(50)$ have been predicted, i.e. all three dynamic variables are predicted for




FIG. 5: Coverage of the prediction confidence intervals for the consecutive model. The horizontal axis is the confidence level $\alpha = \{0.05, 0.1, \dots, 0.95, 0.99\}$ which constitutes the desired coverage of the confidence intervals. The vertical axis is the realized coverage obtained for 100 data realizations. The red error bars are the result obtained for the asymptotic threshold which yield conservative outcomes for predictions of $A(t)$ and $B(t)$. The black error bars indicate the results for the Monte-Carlo thresholds which shows almost perfect agreement with the confidence level in all prediction scenarios.

an early, an intermediate, and a late point in time. For this analysis, a hundred noise realizations have been analyzed. The error bars plotted in this figure are bootstrap confidence intervals of the mean coverage.

The coverage obtained for the asymptotic threshold (red) tends to be conservative, i.e. the true model response is inside the confidence interval more frequently as specified by the confidence level α . This means that there are more false negatives than intended which does not constitute a serious problem in terms of validity of conclusions. In contrast, an anti-conservative coverage would constitute an issue because an increased false positive rate could lead to invalid reasoning.

The coverage obtained by the adjusted thresholds obtained by the Monte-Carlo algorithm shown in Fig. 3 are displayed by the black error bars in Fig. 5. Here, the coverage coincides with the confidence level which confirms the validity respective the prediction profile likelihood based confidence intervals.

4. Comparison of PCI and VCI

The validation profile likelihood satisfies

$$\text{VPL}^{\text{SD}}(z|y) \leq \text{PPL}(z_{\text{alt}}|y) + \frac{1}{2} \frac{(z_{\text{alt}} - z)^2}{\text{SD}^2} \quad \forall z_{\text{alt}} \quad (\text{A6})$$

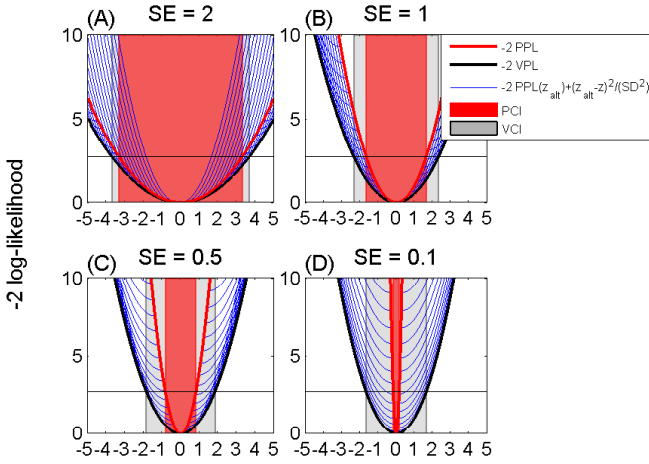


FIG. 6: Comparison of prediction and validation confidence intervals. Panel (A) shows a prediction profile likelihood (red line) with a rather flat shape. Here, the curvature of the prediction profile likelihood corresponds to a prediction standard error $SE = 2$. In this case, the prediction confidence intervals are large (red shaded) and the increase of the validation confidence intervals (gray) is smaller than indicated by the validation data error SD . If the data is more informative, i.e. SE decreases (panels B-D), the slope of prediction profile likelihood increases yielding larger difference between the PCI and VCI.

i.e. the $VPL^{SD}(z|y)$ is smaller than the right hand side of the inequality for any alternative predicted value z_{alt} . On the one hand, the inequality can be utilized to interpret a difference between the respective confidence intervals. Furthermore, the equation can be utilized for consistency checks, e.g. to prove the numerically calculated VPL and PPL. Small differences between the size of the VCI^{SD} and PCI indicate a flat prediction profile likelihood close to the threshold whereas deviations of the confidence intervals in the order of magnitude of SD occur if the PPL has a large slope. This aspect is illustrated in the following.

For illustration purpose, a quadratic prediction profile likelihood

$$-2 PPL(z) = \frac{z^2}{SE^2} \quad (A7)$$

with $SE \in \{0.1, 0.5, 1, 2\}$ has been assumed. These four settings are shown in Fig. 6. The prediction profile likelihood is shown as a red line. For several z_{alt} , the quadratic term in [A6] is plotted by blue curves attached to the PPL. The VPL constitutes the infimum of these curves which in this special case can be calculated analytically and is given by

$$-2 VPL^{SD}(z) = \frac{z^2 SE^2}{(SD^2 + SE^2)^2} + \left(\frac{z}{SD} - \frac{z SE^2}{SD (SD^2 + SE^2)} \right)^2 \quad (A8)$$

Panel (A) shows the comparison for $SE > SD$. In this case, the boundaries of the VCI and the PCI differ only by a value around 0.38. In Panel (B), SE is chosen equal to SD . In Panels (C) and (D) SE is further decreased. This corresponds to more informative data for predicting the exact value of z . In these cases, the optimum of the PPL is narrow in comparison to validation data error SD . Then, during fitting the model, a mismatch $z - z^*$ is predominantly explained by the observation error of the validation data point. The difference of the boundaries of the confidence intervals increase and approach the 10% quantile of the Gaussian distribution, i.e. a value $\text{icdf}(N(0, SD = 1), .95) = 1.64$ which is the one-sided 5% confidence interval for a validation data point for a constant model prediction, i.e. for $SE \rightarrow 0$.

5. Prior information

If prior information about parameters is available, e.g. a prior distribution $\pi(\theta)$, maximum likelihood estimation is replaced by maximum a-posteriori (MAP) estimation

$$\hat{\theta}_{MAP} = \arg \max_{\theta} \rho(y|\theta) \pi(\theta) \quad (A9)$$

$$= \arg \max_{\theta} (\text{LL}(y|\theta) + \log(\pi(\theta))) \quad (A10)$$

i.e. the parameters are estimated by maximizing the a-posterior probability of the data and the parameter estimates. For most common priors, MAP estimation can be performed by MLE using a penalized likelihood. As an example, a log-normal prior for a parameter component θ' yields

$$\text{LL}_{\text{prior}} = \text{LL} - \frac{1}{2} \frac{(\log(\theta') - \langle \log(\theta') \rangle)^2}{\text{Var}(\log(\theta'))} + \text{const} \quad (A11)$$

To incorporate prior knowledge, the presented prediction profile likelihood approach has to be generalized to MAP estimation and the penalized likelihood [A11] is used instead of the standard log-likelihood LL.

To illustrate the incorporation of prior knowledge for parameter values, the initial concentration $A(0) = \theta_3$ is assumed to be drawn from a log-normal distribution

$$\theta_3 \sim \text{log}N(0, 1) \quad (A12)$$

with expectation $\langle \log(\theta_3) \rangle = 0$ and variance $\text{Var}(\log(\theta_3)) = 0.1$. For parameter estimation, this is accounted for by using the penalized likelihood [A11], i.e. by adding an additional term to the residual sum of squares.

As in the example in the main text, the calculation of the prediction and validation confidence intervals has been repeated for $t = 0, 10, \dots, 100$ and all three dynamic states $A(t)$, $B(t)$, $C(t)$. In this example, the true value of $A(0) \equiv \theta_3$ has been drawn according to the prior from the log-normal distribution [A12].

Fig. 7 shows the prediction profile likelihood functions as curves in vertical direction as well as the respective 90% prediction confidence intervals as dark gray shaded regions. The prediction confidence intervals plotted in light color are obtained if θ_3 is estimated without prior information. Because C is the measured compound in our example, the prediction confidence intervals for C are much smaller than for A and B . However, also A and B yield bounded prediction confidence intervals which can be interpreted as *observability* of these dynamic states. Omitting the prior information yields larger prediction confidence intervals, especially for the unobserved states $A(t)$ and $B(t)$.

6. Validation profile likelihood for the consecutive reaction model

In the main text, prediction confidence intervals have been shown for the consecutive reaction model. Fig. 8 shows the corresponding validation profile and the respective validation confidence intervals for the same noise realization for all dynamic variables $A(t)$, $B(t)$, and $C(t)$. Validation confidence intervals account for the measurement noise in a validation experiment. Therefore, they are larger than the prediction confidence intervals shown in the main text in Fig. 1, panels (c)-(e).

Because Gaussian noise $\varepsilon \sim N(\mu, SD^2)$ has been assumed, the validation confidence intervals covers negative values if the true model response $\mu = F(D_{\text{pred}}, \theta_{\text{true}})$ is close to zero.

7. Observability of the long-term dynamics

In the main text, it has been discussed how to exploit the prediction profile likelihood for observability analyses. For the two step model



it has been shown that measurements of the compound C which sample only the transient increase and therefore does not provide information about the steady state level lead to non-observability of $A(t)$, and $B(t)$ for $t > 0$. In addition to that result, Fig. 9 shows prediction confidence intervals for $C(t)$ for times much larger than the measurement times $t = 0, 2, \dots, 20$. $C(t)$ becomes practically non-observable for times which are much larger than the time sampling interval.

8. Characteristics of the MAP kinase model

To demonstrate the applicability of our approach in a realistic setting, the published model of MAP kinase signaling [Kholodenko2000] has been utilized to illustrate

the calculation of prediction and validation confidence intervals.

Fig. 10 shows the long-term dynamics of this model, i.e. the oscillations. In our analysis only the initial phase, i.e. the first 1000 seconds have been considered. This time interval is characterized by strong nonlinearity of the model response with respect to the parameters and constitutes a compromise setting between a transient and an oscillatory dynamics.

Tab. 1 summarizes the model parameters as they have been published in [Kholodenko2000]. The Hill coefficient n is assumed to be equal to one. For the observational noise of the validation data, the same noise level as for the experimental data has been assumed, i.e. $\sigma = SD = 10$.

9. Experimental design conclusions

The size of prediction confidence interval can be utilized to figure out informative experimental designs. If the information of a single data points is intended to be evaluated, then the validation confidence intervals are appropriate. If many experimental replicates are feasible, the average observation will have a small standard error and then prediction confidence intervals can be used to assess a design.

Fig. 11 shows the size of the 90% prediction confidence intervals (upper row), i.e. the difference between the upper and lower boundary, and the size of the validation confidence intervals (lower row) along the time axis. The size is plotted in absolute concentrations (left panels) and relative to the total amount of the protein (right panels).

Independently from the way of the assessment, Erk (blue lines) yields the smallest prediction and validation confidence intervals for $300 < t \leq 1000$. Therefore, measurements of Erk in this time interval constitute very informative experimental designs for testing the model. As discussed in the main text, such a setting which is appropriate for validating the whole model is not informative for improving the model parameters. For such a purpose, new experimental data has to be generated for designs in which the model behavior is not yet precisely specified, i.e. for a setting with large prediction or validation confidence intervals. In these terms, Erk** (gray lines) is most informative in absolute units between 100 and 200 seconds. Also absolute measurements of Mek* (red lines) and Mek** (green lines) along the whole time axis are informative.

If only the amount of a phosphorylated form relative to the total concentration of the protein is experimentally accessible, then the panels on the right should be evaluated to assess the power of a design. In our example, the outcome for the prediction profile likelihood is very similar to the results obtained for absolute concentrations. Again, Erk for $t < 200$ as well as Mek* and Mek** are most informative. For the validation confidence intervals, Raf and Raf* appear as informative because the total concentration of Raf is three times smaller than the

symbol	description	value	lower boundary	upper boundary	units
V_1	max. enzyme rate	2.5	1e-8	1e6	nM s ⁻¹
n	Hill coefficient of the feedback	1	1	1	1
K_I	Michaelis constant	9	1e-8	1e6	nM
K_1	Michaelis constant	10	1e-8	1e6	nM
V_2	max. enzyme rate	0.25	1e-8	1e6	nM s ⁻¹
K_2	Michaelis constant	8	1e-8	1e6	nM
k_3	catalytic rate constant	0.025	1e-8	1e6	nM s ⁻¹
K_3	Michaelis constant	15	1e-8	1e6	nM
k_4	catalytic rate constant	0.025	1e-8	1e6	nM s ⁻¹
K_4	Michaelis constant	15	1e-8	1e6	nM
V_5	max. enzyme rate	0.75	1e-8	1e6	nM s ⁻¹
K_5	Michaelis constant	15	1e-8	1e6	nM
V_6	max. enzyme rate	0.75	1e-8	1e6	nM s ⁻¹
K_6	Michaelis constant	15	1e-8	1e6	nM
k_7	catalytic rate constant	0.025	1e-8	1e6	nM s ⁻¹
K_7	Michaelis constant	15	1e-8	1e6	nM
k_8	catalytic rate constant	0.025	1e-8	1e6	nM s ⁻¹
K_8	Michaelis constant	15	1e-8	1e6	nM
V_9	max. enzyme rate	0.5	1e-8	1e6	nM s ⁻¹
K_9	Michaelis constant	15	1e-8	1e6	nM
V_{10}	max. enzyme rate	0.5	1e-8	1e6	nM s ⁻¹
K_{10}	max. enzyme rate	15	1e-8	1e6	nM

TABLE I: Parameters of the MAP kinase model as published in [Kholodenko2000].

total amount of Mek and Erk.

10. Implementation

The `cvodes` package from the Suite of Nonlinear Differential/Algebraic Equation Solvers (SUNDIALS) [Hindmarsh2005] has been used for the numerical integration of the ODEs and the sensitivity equations. MATLAB's `fmincon` optimizer was used to estimate the parameters. The gradient and the Hessian of the objective function have been provided for the optimizer using the sensitivity equations [Leis1988] and the approximation

$$\frac{\partial^2}{\partial\theta_j\partial\theta_k}\text{LL} \approx \sum_i \frac{1}{\sigma_i^2} \frac{\partial F_i}{\partial\theta_j} \frac{\partial F_i}{\partial\theta_k} \quad (\text{A14})$$

of the second derivatives [Press1992]. Within a single optimization procedure, the parameters have been alter-

nately optimized on the logarithmic scale as well as the common linear scale until the optimizer converged to a common value on both scales.

For the calculation of the validation profile likelihood, 30 test predictions z within the reasonable range of predictions given by the model structure and SD have been evaluated to obtain an initial guess of the likelihood shape. Then the grid is iteratively refined and/or enlarged until a smooth validation likelihood covering the whole confidence interval is obtained and local minima have been removed. For a single profile likelihood, around 10^2 to 10^3 optimizations were required.

For the prediction profile likelihood the initial guess is obtained from the VPL by equation [17] in the main text. The gaps in this guess are then filled by nonlinear constrained optimization. If the constraint optimization procedure did not converge, the validation data error SD has been iteratively decreased by factors $10^0, 10^{-1}, \dots, 10^{-5}$.

[Bayarri2004] M. J. Bayarri and J. O. Berger. The interplay of Bayesian and frequentist analysis. *Statistical Science*, 1:58–80 2004.

[Bjornstad1990] J. F. Bjornstad. Predictive likelihood: A review. *Statistical Science*, 2:242–254 1990.

[Booth1998] J. G. Booth and J. P. Hobert. Standard errors of prediction in generalized linear mixed models. *Journal of the American Statistical Association*, 93:262–267, 1998.

[Butler1986] R. W. Butler. Predictive likelihood inference with applications. *Journal of the Royal Statistical Society, Series B*, 48(1):1–38, 1986.

[Cooley1989] T. F. Cooley, W. R. Parke, and S. Chib. Predictive efficiency for simple nonlinear models. *Journal of Econometrics*, 40(1):33–44, 1989.

[Cox1994] D. Cox and D. V. Hinkley. *Theoretical Statistics*. Chapman & Hall, London, 1994.

[Davison1997] A. C. Davison and D. V. Hinkley. *Bootstrap methods and their application*. Cambridge University Press, Cambridge, 1997.

[DiCiccio1987] T. DiCiccio and R. Tibshirani. Bootstrap confidence intervals and bootstrap approximations. *Journal of the American Statistical Association*, 82(397):163–170,

- 1987.
- [Feder1968] P. I. Feder. On the distribution of the Log Likelihood Ratio test statistic when the true parameter is "near" the boundaries of the hypothesis regions. *The Annals of Mathematical Statistics*, 39(6):2044–2055, 1968.
- [Gelman2003] A. Gelman and J. B. Carlin and H. S. Stern and D. B. Rubin *Bayesian data analysis, Second Edition (Chapman & Hall/CRC Texts in Statistical Science)*. Chapman & Hall, London, 2003.
- [Hindmarsh2005] A. C. Hindmarsh, C. S. Woodward et al. SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Transactions on Mathematical Software*, 31(3):363–396, 2005.
- [Hinkley1979] D. Hinkley. Predictive likelihood. *The Annals of Statistics*, 7(4):718–728, 1979.
- [Hlavacek2009] W. S. Hlavacek. How to deal with large models? *Molecular Systems Biology*, 5:240–242, 2009.
- [Joshi2006] M. Joshi, A. Seidel-Morgenstern, and A. Kremling. Exploiting the bootstrap method for quantifying parameter confidence intervals in dynamical systems. *Metabolic Engineering*, 8(5):447–455, 2006.
- [Kass1998] R. E. Kass, B. P. Carlin, A. Gelman, and R. M. Neal. Markov Chain Monte Carlo in practice: A roundtable discussion. *The American Statistician*, 52(2):93–100, 1998.
- [Kholodenko2000] B. N. Kholodenko. Negative feedback and ultrasensitivity can bring about oscillations in the mitogen-activated protein kinase cascades. *European Journal of Biochemistry*, 267(6):1583–1588, 2000.
- [Kreutz2007] C. Kreutz, J. Timmer et al. An error model for protein quantification. *Bioinformatics*, 23(20):2747–2753, 2007.
- [Kreutz2009] C. Kreutz, and J. Timmer. Systems biology: experimental design. *FEBS Journal*, 276(4):923–942, 2009.
- [Leis1988] J. Leis and M. Kramer. The simultaneous solution and sensitivity analysis of systems described by ordinary differential equations. *ACM Transactions on Mathematical Software*, 14(1):45–60, 1988.
- [Marimont1979] R. B. Marimont and M. B. Shapiro. Nearest Neighbour searches and the curse of dimensionality. *Journal of the Institute of Mathematics and its Applications*, 24: 59–70, 1979.
- [Marjoram2003] P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov Chain Monte Carlo without likelihoods. *PNAS*, 100(26): 15324–15328, 2003.
- [Meeker1995] W. Meeker and L. Escobar. Teaching about approximate confidence regions based on maximum likelihood estimation. *The American Statistician*, 49(1): 48–53, 1995.
- [Molinaro2005] A. M. Molinaro, R. Simon, and R. M. Pfeiffer. Prediction error estimation: A comparison of resampling methods. *Bioinformatics*, 21(15):3301–3307, 2005.
- [Swameye2003] I. Swameye, T. Müller, J. Timmer, O. Sandra, and U. Klingmüller. Identification of nucleocytoplasmic cycling as a remote sensor in cellular signaling by data-based modeling. *PNAS*, 100(3):1028–1033, 2003.
- [Press1992] W. H. Press, B. P. Flannery, S. A. Saul, and W. T. Vetterling. *Numerical recipes*. Cambridge University Press, Cambridge, 1992.
- [Raue2009] A. Raue, U. Klingmüller et al, and J. Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25:1923–1929, 2009.
- [Raue2010] A. Raue, V. Becker, U. Klingmüller, and J. Timmer. Identifiability and observability analysis for experimental design in non-linear dynamical models. *Chaos*, 20(4):045105, 2010.
- [Sachs1984] L. Sachs. *Applied statistics: A handbook of techniques*. Springer, New York, 1984.
- [Scott1991] D. W. Scott, W. David, and M. P. Wand. Feasibility of multivariate density estimates. *Biometrika*, 78:197–205, 1991.
- [Seber1989] G. A. F. Seber and C. J. Wild. *Nonlinear regression*. Wiley, New York, 1989.
- [Venzon1988] D. J. Venzon and S. H. Moolgavkar. A method for computing profile-likelihood-based confidence intervals. *Applied Statistics*, 37(1):87–94, 1988.




FIG. 7: The curves in vertical direction are the prediction profile likelihood functions for $A(t)$ (left panel), $B(t)$ (middle), and $C(t)$ (right panel) if a log-normal prior for θ_3 is assumed. The respective 90% confidence intervals are plotted in dark gray. The light gray regions indicate the 90% confidence intervals if the parameter θ_3 is estimated without prior information.




FIG. 8: The left panel shows the validation confidence intervals for the unobserved state $A(t)$. The validation profile likelihood functions are plotted as curves in vertical direction. For the plotting the confidence intervals along the time axis, the VCIs have been interconnected by cubic piecewise interpolation. Validation confidence intervals and validation profile likelihood functions for the intermediate unobserved state $B(t)$ are shown in the middle and for $C(t)$ in the right panel.




FIG. 9: Prediction confidence intervals for the extrapolation of $C(t)$ to time points much larger than the measurement times. Because in this example the experimental design does not provide sufficient information about the steady state level of $C(t)$, the prediction confidence intervals diverge and the steady state $C(t)$ for $t \rightarrow \infty$ is non-observable.




FIG. 10: The long-term dynamics of the MAP kinase model shows regular oscillations. In our analysis, the first 1000 seconds, as highlighted by gray background color, have been analyzed as a compromise between a transient and oscillatory dynamics.




FIG. 11: Size of the prediction confidence intervals for the dynamic states of the MAP kinase model. The left panels show the size of the confidence intervals in absolute units. In the panels on the right, the size is plotted relative to the total concentration of a protein. The upper row shows the results for the prediction confidence intervals, the lower row for the validation confidence intervals.