Application of Predictive Model Selection to Coupled Models

Gabriel Terejanu , Todd Oliver , Chris Simmons

Abstract—**A predictive Bayesian model selection approach is presented to discriminate coupled models used to predict an unobserved quantity of interest (QoI). The need for accurate predictions arises in a variety of critical applications such as climate, aerospace and defense. A model problem is introduced to study the prediction yielded by the coupling of two physics/sub-components. For each single physics domain, a set of model classes and a set of sensor observations are available. A goal-oriented algorithm using a predictive approach to Bayesian model selection is then used to select the combination of single physics models that best predict the QoI. It is shown that the best coupled model for prediction is the one that provides the most robust predictive distribution for the QoI.**

Index Terms—**Predictive Model Selection, Quantity of Interest, Model Validation, Decision Making, Bayesian Analysis**

I. INTRODUCTION

WITH the exponential growth of available computing power and the continued development of advanced numerical algorithms, computational science has undergone a revolution in which computer models are used to simulate increasingly complex phenomena. Additionally, such simulations are guiding critical decisions that affect our welfare and security, such as climate change, performance of energy and defense systems and the biology of diseases. Reliable predictions of such complex physical systems requires sophisticated mathematical models of the physical phenomena involved. But also required is a systematic, comprehensive treatment of the calibration and validation of the models, as well as the quantification of the uncertainties inherent in such models.

While recently some attention has been paid to the propagation of uncertainty, considerably less attention has been paid to the validation of these complex, multiphysics models. This becomes particularly challenging when the quantity of interest (QoI) cannot be directly measured, and the comparison of model predictions with real data is not possible. Such QoIs may be the catastrophic failure of the thermal protection system of a space shuttle reentering the atmosphere or the different performance characteristics of nuclear weapons in order to maintain the nuclear stockpile without undergoing underground nuclear testing.

In this paper, we present an intuitive interpretation of the predictive model selection in the context of Bayesian analysis. While the predictive model selection is not an new

idea, see Refs.[1], [2], [3], here we emphasize the connection between the QoI-aware evidence and the Bayesian model averaging used for estimation. This new interpretation of the Bayesian predictive model selection reveals that the best model for prediction is the one which provides the most robust predictive probability density function (pdf) for the QoI.

The latest advances in Markov Chain Monte Carlo (MCMC) [4] and estimators based on the k-nearest neighbor [5] are used to compute the information theoretic measures required in the problem of predictive selection of coupled models. It is further argued that equivalence between predictive model selection and conventional Bayesian model selection can be reached by performing optimal experimental design [6] for model discrimination.

The structure of the paper is as follows: first the selection problem of coupled models is stated in Section II. The conventional Bayesian model selection is described in Section III and the extension to QoI-aware evidence is derived in Section IV. The model problem and numerical results are presented in Section V and Section VI respectively. The conclusions and future work are discussed in Section VII.

II. PROBLEM STATEMENT

Here we are interested in the prediction of a coupled model. The problem of selecting the best coupled model in the context of the QoI is to find the combination of single physics models that best predict an unobserved QoI in some sense, see Fig. 1. Thus at the single physics level, we have two physics A and B, each with a model class set, $\mathcal{M}^A$ and $\mathcal{M}^B$, and a set of observations D^A and D^B respectively. The cardinality of the two sets of model classes are $|\mathcal{M}^A| = K_A$, $|\mathcal{M}^B| = K_B$, and the the definition of model classes in each set is given by the state equations and measurement models as follows,

$$M_i^A : \begin{cases} \mathbf{r}_i^A(\mathbf{u}_i^A, \boldsymbol{\theta}_i^A) = \mathbf{0} \\ \mathbf{y}^A = \mathbf{y}_i^A(\mathbf{u}_i^A, \boldsymbol{\theta}_i^A) \end{cases} \qquad M_j^B : \begin{cases} \mathbf{r}_j^B(\mathbf{u}_j^B, \boldsymbol{\theta}_j^B) = \mathbf{0} \\ \mathbf{y}^B = \mathbf{y}_j^B(\mathbf{u}_j^B, \boldsymbol{\theta}_j^B). \end{cases}$$

All the possible couplings of single physics models yield the set of coupled model classes $\mathcal{M} = \{M_{ij} = (M_i^A, M_j^B) | M_i^A \in \mathcal{M}^A, M_j^B \in \mathcal{M}^B\}$ with cardinality $|\mathcal{M}| = K_A K_B$. The definition of a coupled model class in this set is given by the state and measurement equation, and in addition we also have the model for the QoI,

$$M_{ij} : \begin{cases} \mathbf{r}(\mathbf{u}_i^A, \mathbf{u}_j^B, \boldsymbol{\theta}_i^A, \boldsymbol{\theta}_j^B) = \mathbf{0} \\ \mathbf{y}^{AB} = \mathbf{y}_{ij}^{AB}(\mathbf{u}_i^A, \mathbf{u}_j^B, \boldsymbol{\theta}_i^A, \boldsymbol{\theta}_j^B) \\ \mathbf{q} = \mathbf{q}_{ij}^{AB}(\mathbf{u}_i^A, \mathbf{u}_j^B, \boldsymbol{\theta}_i^A, \boldsymbol{\theta}_j^B). \end{cases}$$
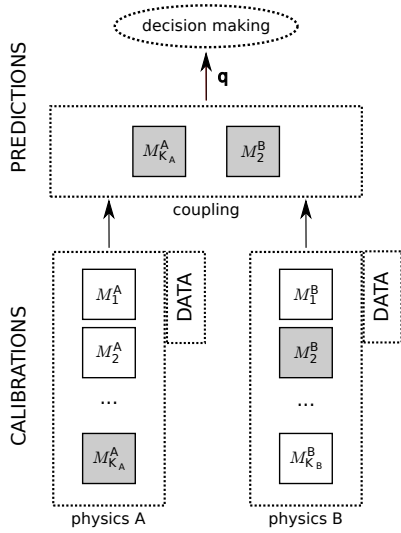
Fig. 1.   Predictive Selection for Coupled Models

Having the set of all the coupled models, the selection problem becomes finding the best coupled model in the set, $\mathcal{M}$, for prediction purposes.

III. BAYESIAN MODEL SELECTION

In the context of $\mathcal{M}$-closed perspective, the conventional Bayesian approach to model selection is to choose the model which has the highest posterior plausibility,

$$M^* = \arg\max_M \pi(M|D, \mathcal{M}). \quad (1)$$

Given the data at the single physics level, one can compute the posterior model plausibility for all the models in the $\mathcal{M}^A$ and $\mathcal{M}^B$ sets, as the product between evidence and prior plausibility,

$$\pi(M_i^A|D^A, \mathcal{M}) \propto \pi(D^A|M_i^A, \mathcal{M})\pi(M_i^A|\mathcal{M}). \quad (2)$$

The evidence is obtained during the calibration process for each single-physics models, and it is given by the normalization constant in the Bayes rule, used to compute the posterior pdf of model parameters,

$$\pi(\boldsymbol{\theta}_i^A|D^A, M_i^A, \mathcal{M}) = \frac{\pi(D^A|\boldsymbol{\theta}_i^A, M_i^A, \mathcal{M})\pi(\boldsymbol{\theta}_i^A|M_i^A, \mathcal{M})}{\pi(D^A|M_i^A, \mathcal{M})}.$$

With no data at the coupled level, one can easily obtained the posterior plausibility for coupled models as,

$$\pi(M_{ij}|D, \mathcal{M}) = \pi(M_i^A|D^A, \mathcal{M})\pi(M_j^B|D^B, \mathcal{M}). \quad (3)$$

Notice, that the best coupled model is given by coupling the best models at the single physics level. Computationally this is advantageous as there is no need to build all the possible coupled models using the single physics models. Looking at just two coupled models, we can say that we prefer model M_1 over model M_2 if and only if $\pi(M_1|D, \mathcal{M}) > \pi(M_2|D, \mathcal{M})$. This inequality can be recasted as the following product of ratios:

$$\underbrace{\frac{\pi(D|M_1)}{\pi(D|M_2)}}_{\text{Bayes factor}} \underbrace{\frac{\pi(M_1, \mathcal{M})}{\pi(M_2, \mathcal{M})}}_{\text{prior odds}} > 1 \quad (4)$$

The log-evidence is the trade-off between model complexity and how well the model fits the data. In other words the evidence yields the best model that obeys the law of parsimony,

$$\ln[\pi(D|M_1)] = E\left[\ln[\pi(D|\boldsymbol{\theta}, M_1)]\right]$$
$$- \mathrm{KL}\Big(\pi(\boldsymbol{\theta}|D, M_1) \,||\, \pi(\boldsymbol{\theta}|M_1)\Big), \quad (5)$$

where the model complexity is given by the Kullback-Leibler divergence between posterior pdf and prior pdf of model parameters [4]. Therefore, this model selection scheme makes use of the following information in choosing the best model:

(model complexity, data fit, prior knowledge) .

Since it obeys Occam's razor, we implicitly gain some robustness with respect to predictions. However, if we have two different QoIs that we would like to predict, it is not obvious if the model selected under this scheme will be able to provide equally good predictions for both QoIs. This is due to the fact that the information about the QoI is not explicitly used in the selection criterion. In the following sections, we will present an extension of model selection scheme to also account for the QoI, and discuss its implications.

IV. PREDICTIVE MODEL SELECTION

Given a model class set $\mathcal{M} = \{M_1, M_2, \ldots, M_K\}$ (for simplicity the double index will be ignored in the model class notation), and a set of observations $D = \{\mathbf{d}_1, \mathbf{d}_2, \ldots, \mathbf{d}_n\}$ generated by an unknown model from $\mathcal{M}$, we are concerned with the problem of selecting the best model to predict an unobserved quantity of interest $\mathbf{q}$. The selection of the model which best estimates the pdf of the QoI is seen here as a decision problem [3]. First one has to find the predictive distribution for each model class and the selection of the best model class is based on the utility of its predictive distribution. Given all the available information, the Bayesian predictive distribution conditioned on a model M_j is given by:

$$\pi(\mathbf{q}|D, M_j) = \int \pi(\mathbf{q}|\boldsymbol{\theta_j}, D, M_j) \,\pi(\boldsymbol{\theta_j}|D, M_j) \,d\boldsymbol{\theta_j} \quad (6)$$

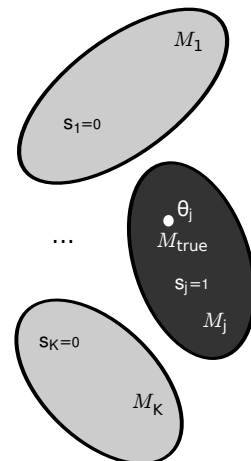where the posterior pdf for model parameters is computed using Bayes rule.

Fig. 2.   Assumption regarding the true model

In the followings it is assumed that the true distribution of the QoI is generated by a model $M_j(\boldsymbol{\theta}_j) \in M_j$, called the true model, see Fig.2. Thus, the true pdf of the QoI can be written as:

$$\pi(\mathbf{q}|\boldsymbol{\theta}, \mathbf{s}, D, \mathcal{M}) = \sum_{j=1}^{K} \pi(\mathbf{q}|\boldsymbol{\theta}_j, D, M_j)s_j \qquad (7)$$

where $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_K)$, $\mathbf{s} = (s_1, s_2, \ldots, s_K)$, and $s_j = 1$ if and only if the true model belongs to model class M_j. Let $U_{\mathbf{q}}(\boldsymbol{\theta}, \mathbf{s}, M_j)$ describe the utility of choosing the pdf associated with model class M_j as the predictive distribution of the QoI, $\mathbf{q}$. Here the utility function is defined as the negative Kullback-Leibler divergence of the true distribution and the predictive distribution of M_j:

$$U_{\mathbf{q}}(\boldsymbol{\theta}, \mathbf{s}, M_j) = -\text{KL}\Big(\pi(\mathbf{q}|\boldsymbol{\theta}, \mathbf{s}, D, \mathcal{M}) \,||\, \pi(\mathbf{q}|D, M_j)\Big) \quad (8)$$

The model that maximizes the following expected utility (QoI-aware evidence) is the model of choice for predictive purposes:

$$M^* = \arg\max_{M_j \in \mathcal{M}} \underbrace{\int U_{\mathbf{q}}(\boldsymbol{\theta}, \mathbf{s}, M_j)\pi(\boldsymbol{\theta}, \mathbf{s}|D, \mathcal{M})d\boldsymbol{\theta}d\mathbf{s}}_{E_{\boldsymbol{\theta}, \mathbf{s}}[U_{\mathbf{q}}(\boldsymbol{\theta}, \mathbf{s}, M_j)]} \quad (9)$$

With few mathematical manipulations, the expected utility can be written as follows,

$$E_{\boldsymbol{\theta}, \mathbf{s}}[U_{\mathbf{q}}(\boldsymbol{\theta}, \mathbf{s}, M_j)] = \int U_{\mathbf{q}}(\boldsymbol{\theta}, \mathbf{s}, M_j)\pi(\boldsymbol{\theta}|\mathbf{s}, D, \mathcal{M})\pi(\mathbf{s}|D, \mathcal{M})d\boldsymbol{\theta}d\mathbf{s}$$

$$= \sum_{i=1}^{K} \pi(M_i|D, \mathcal{M})\int U_{\mathbf{q}}(\boldsymbol{\theta}_i, s_i, M_j)\pi(\boldsymbol{\theta}_i|D, M_i)d\boldsymbol{\theta}_i \qquad (10)$$

$$= -\sum_{i=1}^{K} \pi(M_i|D, \mathcal{M})E_{\boldsymbol{\theta}_i}\Big[\text{KL}\Big(\pi(\mathbf{q}|\boldsymbol{\theta}_i, D, M_i)||\pi(\mathbf{q}|D, M_j)\Big)\Big]$$

A. Interpretation of the expected utility for model selection

Consider now that only two model classes exist in our model set $\mathcal{M} = \{M_1, M_2\}$. We prefer model class M_1 over M_2 and write $M_1 \succ M_2$ if and only if:

$$E_{\boldsymbol{\theta}, \mathbf{s}}[U_{\mathbf{q}}(\boldsymbol{\theta}, \mathbf{s}, M_1)] > E_{\boldsymbol{\theta}, \mathbf{s}}[U_{\mathbf{q}}(\boldsymbol{\theta}, \mathbf{s}, M_2)] \qquad (11)$$

Substituting Eq.(10) into Eq.(11) the following model selection criterion can be derived:

$$\underbrace{\frac{R(M_1||M_2)}{R(M_2||M_1)}}_{\text{Risk ratio}} \underbrace{\frac{\pi(D|M_1)}{\pi(D|M_2)}}_{\text{Bayes factor}} \underbrace{\frac{\pi(M_1, \mathcal{M})}{\pi(M_2, \mathcal{M})}}_{\text{Prior odds}} > 1 \qquad (12)$$

where the numerator in the predictive risk ratio is given by the following expressions. The denominator is obtained by analogy with the numerator.

$$R(M_1||M_2) = E_{\boldsymbol{\theta}_1}\Big[\text{KL}\Big(\pi(\mathbf{q}|\boldsymbol{\theta}_1, D, M_1) \,||\, \pi(\mathbf{q}|D, M_2)\Big)\Big]$$
$$- E_{\boldsymbol{\theta}_1}\Big[\text{KL}\Big(\pi(\mathbf{q}|\boldsymbol{\theta}_1, D, M_1) \,||\, \pi(\mathbf{q}|D, M_1)\Big)\Big] \quad (13)$$

The model selection criterion in Eq.(12) can be interpreted as the evidence of model class M_1 in favor of model class M_2, and is composed of prior evidence given by the prior odds, experimental evidence given by the Bayes factor and the predictive risk ratio which accounts for the loss of choosing the wrong model. According to Trottini and Spezzaferri [3], the expectations in the above ratio have the following meaning: $E_{\boldsymbol{\theta}_1}\Big[\text{KL}\Big(\pi(\mathbf{q}|\boldsymbol{\theta}_1, D, M_1) \,||\, \pi(\mathbf{q}|D, M_2)\Big)\Big]$ - the risk of choosing model class M_2 when the true model belongs to M_1; $E_{\boldsymbol{\theta}_1}\Big[\text{KL}\Big(\pi(\mathbf{q}|\boldsymbol{\theta}_1, D, M_1) \,||\, \pi(\mathbf{q}|D, M_1)\Big)\Big]$ - even if we report the distribution $\pi(\mathbf{q}|D, M_1)$ when the true model belongs to M_1, there is a risk incurred due to the unknown value of $\boldsymbol{\theta}_1$ that generated the true model. Comparing with the previous model selection scheme, the following information is used in this scheme to select the best model:

(**QoI**, model complexity, data fit, prior knowledge) .

B. Calculation of the expected utility used for predictive model selection

The calculation of the QoI-aware evidence in Eq.(10) is challenging as we are dealing with high dimensional integrals, and the number of samples in the posterior distributions is dependent on the MCMC algorithms and computational complexity of the forward model. Thus, we would like to simplify this calculation. Starting from Eq.(10) the following expression for the expected utility can be obtained:

$$E_{\boldsymbol{\theta}, \mathbf{s}}[U_{\mathbf{q}}(\boldsymbol{\theta}, \mathbf{s}, M_j)] = -\sum_{i=1}^{K} \pi(M_i|D, \mathcal{M})\int \pi(\boldsymbol{\theta}_i|D, M_i)\pi(\mathbf{q}|\boldsymbol{\theta}_i, D, M_i)\log\frac{\pi(\mathbf{q}|\boldsymbol{\theta}_i, D, M_i)}{\pi(\mathbf{q}|D, M_j)}d\boldsymbol{\theta}_id\mathbf{q}$$

$$= -\sum_{i=1}^{K} \pi(M_i|D, \mathcal{M})\int \pi(\mathbf{q}, \boldsymbol{\theta}_i|D, M_i)\log\pi(\mathbf{q}|\boldsymbol{\theta}_i, D, M_i)d\boldsymbol{\theta}_id\mathbf{q} + \int \pi(\mathbf{q}|D, \mathcal{M})\log\pi(\mathbf{q}|D, M_j)d\mathbf{q} \quad (14)$$

Where the predictive pdf under all models is given by,

$$\pi(\mathbf{q}|D, \mathcal{M}) = \sum_{i=1}^{K} \pi(M_i|D, \mathcal{M})\pi(\mathbf{q}|D, M_i). \qquad (15)$$

Since the first term in Eq.(14) is the same for all models M_j, for $j = 1 \ldots K$, the optimization in Eq.(9) is equivalent with maximizing the second term in Eq.(14), which is the negative cross-entropy between the predictive distribution conditioned on all the models and the predictive distribution conditioned on the jth model:

$$M^* = \arg\max_{M_j \in \mathcal{M}} -\text{H}\Big(\pi(\mathbf{q}|D, \mathcal{M}), \pi(\mathbf{q}|D, M_j)\Big) \qquad (16)$$

By writing the optimization as a minimization instead of a maximization and subtracting the entropy of the predictive distribution conditioned on all the models, then one can rewrite the model selection problem as,

$$M^* = \arg \min_{M_j \in \mathcal{M}} \mathrm{H}\Big(\pi(\mathbf{q}|D, \mathcal{M}), \pi(\mathbf{q}|D, M_j)\Big) - \mathrm{H}\Big(\pi(\mathbf{q}|D, \mathcal{M})\Big)$$

$$= \arg \min_{M_j \in \mathcal{M}} \mathrm{KL}\Big(\pi(\mathbf{q}|D, \mathcal{M})\,\Big\|\,\pi(\mathbf{q}|D, M_j)\Big) \qquad (17)$$

Thus, the best model to predict an unobserved quantity of interest $\mathbf{q}$ is the one whose predictive distribution best approximates the predictive distribution conditioned on all the models. This is rather intuitive as all we can say about the unobserved quantity of interest is encoded in the predictive distribution conditioned on all the models. The predictive pdf under all models being the most robust estimate of the QoI for this problem.

This model selection scheme reveals that when the posterior model plausibility is not able to discriminate between the models, the prediction obtained with the selected model is the most robust prediction we can obtain with one model. Thus we are able to account for model uncertainty when predicting the QoI. On the other hand, in the limit, for discriminatory observations, when the posterior plausibility is one for one of the models, the two selection schemes become equivalent.

V. MODEL PROBLEM

The model problem consists of a spring-mass-damper system that is driven by an external force. The spring-mass-damper and the forcing function are considered to be separate physics such that the full system model consists of a coupling of the dynamical system modeling the spring-mass-damper system and a function modeling the forcing. In this model problem, synthetic data are generated according to a truth system.

A. Models

This section describes the models that will form the sets of interest for the single physics. The models of the spring-mass-damper system take the following form:

$$m\ddot{x} + c\dot{x} + \tilde{k}(x)x = 0. \qquad (18)$$

The mass is assumed to be perfectly known, $m = 1$, and the damping coefficient c is a calibration parameter. Model form uncertainty is introduced through the spring models $\tilde{k}(x)$. Three models are considered: a linear spring (**OLS**), a cubic spring (**OCS**), and a quintic spring (**OQS**), given by the following relations:

$$\tilde{k}_{OLS}(x) = k_{1,0} \qquad (19)$$
$$\tilde{k}_{OCS}(x) = k_{3,0} + k_{3,2}x^2 \qquad (20)$$
$$\tilde{k}_{OQS}(x) = k_{5,0} + k_{5,2}x^2 + k_{5,4}x^4 \qquad (21)$$

The models of the forcing function are denoted $\tilde{f}(t)$. Three models are considered: simple exponential decay (**SED**), oscillatory linear decay (**OLD**), and oscillatory exponential decay (**OED**):

$$\tilde{f}_{SED}(t) = F_0 \exp(-t/\tau) \qquad (22)$$

$$\tilde{f}_{OLD}(t) = \begin{cases} F_0(1 - t/\tau)\left[\alpha \sin(\omega t) + 1\right], \ 0 \le t \le \tau \\ \qquad\qquad 0, \quad t > \tau \end{cases}$$
$$\qquad (23)$$

$$\tilde{f}_{OED}(t) = F_0 \exp(-t/\tau)\left[\alpha \sin(\omega t) + 1\right] \qquad (24)$$

The coupling of the spring-mass-damper and the forcing is trivial. Thus, only a single coupling model is considered, and the coupled model is given by

$$m\ddot{x} + c\dot{x} + \tilde{k}(x)x = \tilde{f}(t). \qquad (25)$$

There are three choices for $\tilde{k}(x)$ and three choices for $\tilde{f}(t)$, leading to nine total coupled models.

B. The Truth System

To evaluate the two selection schemes: Bayesian model selection, and predictive model selection, we can construct the true system which will be used to generate data and give the true value of the QoI. The comparison of the two selection criteria will be done with respect to different subsets of models, and the ability of the best model to predict the true value of the QoI. The true model is described by **OQS-OED**:

$$m\ddot{x} + c\dot{x} + \tilde{k}_{OQS}(x)x = \tilde{f}_{OED} \qquad (26)$$

where $m = 1$, $c = 0.1$. The parameters for the spring model are set to $k_{5,0} = 4$, $k_{5,2} = -5$, and $k_{5,4} = 1$. The true forcing function is given by the following values for the parameters: $F_0 = 1$, $\tau = 2\pi$, $\alpha = 0.2$, $\omega = 2$.

The QoI of the coupled model is assumed to be the maximum velocity $\dot{x}_{max} = \max_{t \in \mathbb{R}^+} |\dot{x}(t)|$. The observable for physics A is given by the kinetic energy versus time: $\frac{1}{2}\dot{x}(t_i)^2$ for $i = 1, \ldots, N$. Note that this contains the same information as the velocity except that it is ambiguous with respect to the sign. The observable for physics B is given by the force versus time: $f(t_i)$ for $i = 1, \ldots, M$. In both cases simulated observations have been generated by perturbing the deterministic predictions of the true model, with a log-normal multiplicative noise with standard deviation of 0.1.

VI. NUMERICAL RESULTS

The inverse problem of calibrating the model parameters from the measurement data is solved using MCMC simulations. In our simulations, samples from the posterior distribution are obtained using the statistical library QUESO [7], [4] equipped with the Hybrid Gibbs Transitional Markov Chain Monte Carlo method proposed in Ref. [8]. One advantage of this MCMC algorithm is that it provides an accurate estimate of the log-evidence using the adaptive thermodynamic integration. Estimators based on k-nearest neighbor are used to compute the Kullback-Leibler divergence in Eq.(17), see Appendix. The use of these estimators is advantageous especially when only samples are available to describe the underlying distributions.

Three different scenarios are constructed to assess the predictive capability of the models selected using the two selections schemes: Bayesian model selection and predictive model selection. All the uncertain parameters of the models are considered uniformly distributed and the model error has also been calibrated and propagated to the QoI.

Case 1. First, all the models are included in the two sets, including the components used to generate the true model. For oscillators the model class set is given by $\mathcal{M}^A = \{M_{OLS}, M_{OCS}, M_{OQS}\}$ and for forcing function $\mathcal{M}^B = \{M_{SED}, M_{OLD}, M_{OED}\}$. A number of 10 measurements have been considered for the oscillators and 61

(a) Case 1: QoI pdf

(b) Case 2: QoI pdfs

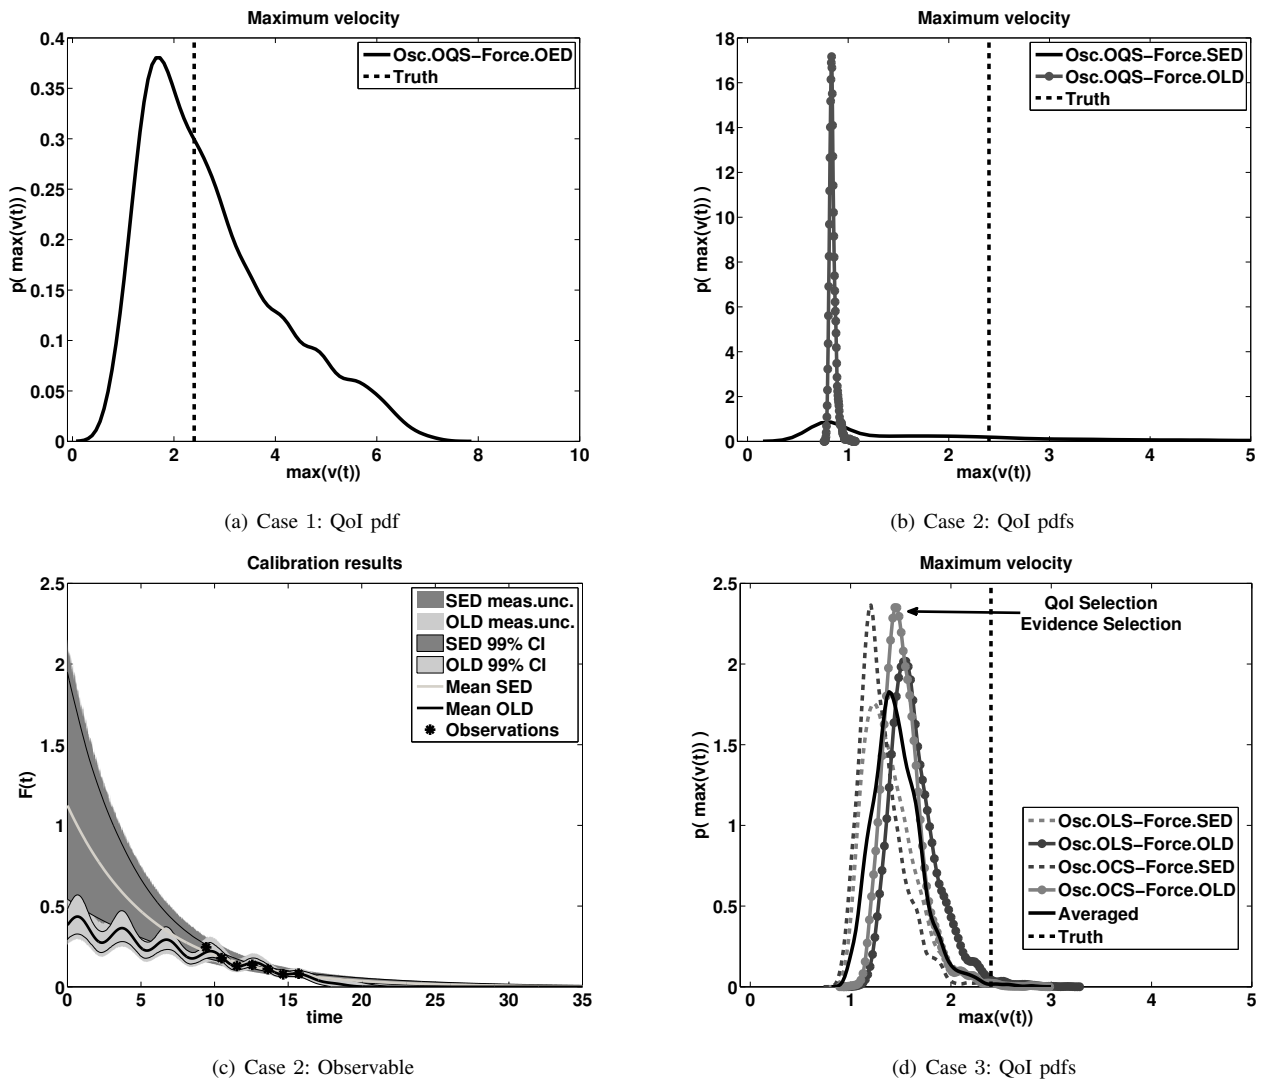(c) Case 2: Observable

(d) Case 3: QoI pdfs

Fig. 3.   Results for the three cases considered

for the forcing. Table I summarizes the results obtained after applying the two approaches. On the first column and the first row, under each model one can find the model plausibility after calibration, and the first number in a cell gives the plausibility of the corresponding coupled model. The number in the parenthesis is the KL divergence used in the predictive model selection.

In this case, the observations provided are enough to discriminate the models at single physics level, oscillator **OQS** and the forcing **OED** are selected in this case. Here, the predictive model selection is consistent with the plausibility based model selection. Notice that the true model belongs to the model class **OQS-OED**, and the prediction of the selected model covers the true value of the QoI, see Fig.3a.

One computational advantage is that when discriminatory observations are available, one does not need to carry the analysis on all the coupled models, just on the coupling of the best single physics ones and still be in agreement with predictive requirements. We argue that for very complex and hierarchical systems, with multiple levels of coupling, such a situation should be preferred and exploited by designing experiments to collect measurements intended to discriminate models.

TABLE I
RESULTS CASE 1

	SED	**OLD**	**OED***
	0.00	0.00	**1.00**
OLS	0.00	0.00	0.00
0.00	(3.75)	(3.36)	(3.67)
OCS	0.00	0.00	0.00
0.00	(4.19)	(3.91)	(4.24)
OQS*	0.00	0.00	**1.00**
1.00	(1.15)	(5.09)	(**0.00**)

Case 2. For the second scenario we remove the forcing that generated the true model. Now the sets of model classes are given by $\mathcal{M}^A = \{M_{OLS}, M_{OCS}, M_{OQS}\}$ and $\mathcal{M}^B = \{M_{SED}, M_{OLD}\}$. The same number of observations are considered for the oscillators and 7 measurements for the forcing models. The results are presented in Table II. As before, we are able to discriminate the oscillators, however we cannot say the same for the forcing models. This can be seen in Fig. 3c, where we can see the prediction of the observable with the two forcing models after calibration. The data supports almost equally well both forcing functions.

In this case the two selection schemes yield two different models: **OQS-OLD** for the Bayesian model selection and **OQS-SED** for the predictive model selection, see Table II.

Looking at their predictions for the QoI, Fig. 3b, we see that the pdf provided by the model selected using conventional Bayesian model selection doesn't even cover the true value of the QoI, whereas the one chosen by the predictive selection scheme covers the true value of the QoI in the tail.

Thus, while the plausibility based selection ignores the model uncertainty, the predictive selection approach yields the model with the most robust predictive pdf for the QoI. This prediction incorporates as much as possible model uncertainty that one can obtain with just one model. Therefore, the predictive approach is recommended in the case when discriminatory observations are not available and one model has to be chosen instead of model averaging, especially for complex hierarchical systems.

TABLE II
RESULTS CASE 2

	SED	OLD
	0.44	**0.56**
OLS	0.00	0.00
0.00	(6.77)	(7.35)
OCS	0.00	0.00
0.00	(5.66)	(7.23)
OQS*	0.44	**0.56**
1.00	**(0.62)**	(1.64)

TABLE III
RESULTS CASE 3

	SED	OLD
	0.48	**0.52**
OLS	0.05	0.06
0.12	(0.39)	(0.82)
OCS	0.42	**0.45**
0.88	(0.37)	**(0.34)**

Case 3. Lastly, we are not including in the model sets any of the components that generated the true model. Thus, $\mathcal{M}^A = \{M_{OLS}, M_{OCS}\}$ and $\mathcal{M}^B = \{M_{SED}, M_{OLD}\}$. In this case only 5 observations are considered for the oscillators and 4 for the forcing functions. We can see from Table III that at the single physics level we are not able to discriminate the oscillators or the forcing functions. For the coupled models both approaches choose the same model **OCS-OLD**, however looking at the predictions of all coupled models, including their average prediction we see that all of them give very low likelihood for the true value of the QoI.

In this case we emphasize that the prediction approach has the same drawback as the Bayesian model selection and Bayesian model averaging. Mainly, we are at the mercy of our hypotheses and the only way to escape this case is to generate additional hypotheses.

VII. CONCLUSIONS

In this paper, we have presented a model selection criterion that accounts for the predictive capability of coupled models. This is especially useful when complex/multiphysics models are used to calculate a quantity of interest. It has been shown that the prediction obtain with the model chosen by the predictive approach, is the most robust prediction that one can obtain with just one model. This is because in part it incorporates model uncertainty, while conventional Bayesian model selection ignores it. For discriminatory measurements the Bayesian model selection and predictive model selection are equivalent. This suggests that when additional data collection is possible then designing experiments for model discrimination is computationally preferred for complex models.

APPENDIX

The approximation of the Kullback-Leibler divergence is based on a k-nearest neighbor approach [5].

$$\text{KL}\left(p(x|D_n)\middle\|p(x|D_{n-1})\right) \approx \frac{d_X}{N_n}\sum_{i=1}^{N_n}\log\frac{\nu_{n-1}(i)}{\rho_n(i)} + \log\frac{N_{n-1}}{N_n-1}$$

where d_X is the dimensionality of the random variable X, N_n and N_{n-1} give the number of samples $\{X_n^i | i = 1, \ldots, N_n\} \sim p(x|D_n)$ and $\{X_{n-1}^j | j = 1, \ldots, N_{n-1}\} \sim p(x|D_{n-1})$ respectively, and the two distances $\nu_{n-1}(i)$ and $\rho_n(i)$ are defined as follows:

$$\rho_n(i) = \min_{j=1\ldots N_n, j\neq i}\|X_n^i - X_n^j\|_\infty$$
$$\nu_{n-1}(i) = \min_{j=1\ldots N_{n-1}}\|X_n^i - X_{n-1}^j\|_\infty$$

REFERENCES

[1] S. Geisser and W. Eddy, "A predictive approach to model selection," *Journal of the American Statistical Association*, vol. 74(365), pp. 153–160, 1979.
[2] A. S. Martini and F. Spezzaferri, "A predictive model selection criterion," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 46(2), pp. 296–303, 1984.
[3] M. Trottini and F. Spezzaferri, "A generalized predictive criterion for model selection," *Canadian Journal of Statistics*, vol. 30, pp. 79–96, 2002.
[4] S. H. Cheung and J. L. Beck, "New bayesian updating methodology for model validation and robust predictions of a target system based on hierarchical subsystem tests." *Computer Methods in Applied Mechanics and Engineering*, 2009, accepted for publication.
[5] Q. Wang, S. Kulkarni, and S. Verdu, "A nearest-neighbor approach to estimating divergence between continuous random vectors," in *Information Theory, 2006 IEEE International Symposium on*, 2006, pp. 242–246.
[6] G. Terejanu, R. R. Upadhyay, and K. Miki, "Bayesian Experimental Design for the Active Nitridation of Graphite by Atomic Nitrogen," *Experimental Thermal and Fluid Science, under review, arXiv:1107.1445v1*, 2011.
[7] E. E. Prudencio and K. W. Schulz, "The Parallel C++ Statistical Library 'QUESO': Quantification of Uncertainty for Estimation, Simulation and Optimization," *Submitted to IEEE IPDPS*, 2011.
[8] S. H. Cheung and J. L. Beck, "New bayesian updating methodology for model validation and robust predictions based on data from hierarchical subsystem tests," California Institute of Technology, Tech. Rep. EERL 2008-04, 2008.