

MIT Open Access Articles

A general model validation and testing tool

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Vanslette, Kevin et al., "A general model validation and testing tool." Reliability Engineering & System Safety 195 (March 2020): 106684 doi. 10.1016/j.ress.2019.106684 ©2019

As Published: <https://dx.doi.org/10.1016/J.RESS.2019.106684>

Publisher: Elsevier BV

Persistent URL: <https://hdl.handle.net/1721.1/127676>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs License



A General Model Validation and Testing Tool

Kevin Vanslette, Tony Tohme, and Kamal Youcef-Toumi

Department of Mechanical Engineering
Massachusetts Institute of Technology, Cambridge MA 02139, USA

October 28, 2019

Abstract

We construct and propose the “Bayesian Validation Metric” (BVM) as a general model validation and testing tool. We find the BVM to be capable of representing all of the standard validation metrics (square error, reliability, probability of agreement, frequentist, area, probability density comparison, statistical hypothesis testing, and Bayesian model testing) as special cases and find that it can be used to improve, generalize, or further quantify their uncertainties. Thus, the BVM allows us to assess the similarities and differences between existing validation metrics in a new light.

The BVM has the capacity to allow users to invent and select models according to novel validation requirements. We formulate and test a few novel compound validation metrics that improve upon other validation metrics in the literature. Further, we construct the BVM Ratio for the purpose of quantifying model selection under user defined definitions of agreement in the presence or absence of uncertainty. This construction generalizes the Bayesian model testing framework.

1 Introduction

Pivotal to the scientific and engineering process is the testing of hypothesis and the validation of mathematical and computational models. Without a validation procedure, there is no reason to believe that a model, which has been designed to serve as a convenient representation of physical and/or human-made processes, has fulfilled its functional purpose. Thus, validation is the result of the positive justification of a model’s representation of relevant features in the real world [1, 2, 3].

We are interested in studying the validation of multivariate computational models that represent uncertain situations and/or data. It is understood that complete certainty is a special case of uncertainty. The uncertainty in a model or data set may originate from stochasticity, model parameter and input data uncertainty, measurement uncertainty, or other possible aleatoric or epistemic sources of uncertainty. Each of the following data modeling schemes may include quantifiable amounts of uncertainty (or certainty) that we would like to validate on the basis of a set of validation data: neural networks and AI models, machine learning models, Gaussian Process Regression models [4], polynomial chaos and other surrogate models [5, 6, 7, 8], spatial and time series stochastic models, physics based models (usually solutions to differential equations), engineering based models (which are sufficiently abstracted physics based models), Monte Carlo simulation models [9, 10], and more. Model output uncertainties may be quantified through uncertainty propagation techniques (that may or may not include verification, calibration, and validation) [4, 6, 7, 8, 11, 12, 13, 14, 16, 21, 22].

There exist several validation metrics. Each metric is designed to compare features of a model-data pair to quantify validation: square error compares the difference in the data and model values in a point to point or interval fashion [23], the reliability metric [24] and the probability of agreement [25] compare continuous model outputs and data expectation values (the model reliability metric was extended past expectation values in [26]), the frequentist validation metric [27, 28] and statistical hypothesis testing compare data and model test statistics, the area metric compares the cumulative distribution of the model to the estimated cumulative

distribution of the data [15, 16, 17, 18, 19, 20], probability density function (pdf) comparison metrics such as the KL Divergence that represent “closeness” between pdfs, and Bayesian model testing compares the posterior probability that each model would correctly output the observed data [12, 13, 29, 30, 31, 32, 33]. A detailed review of the majority of these metrics may be found in [34, 35, 36, 37] and the references therein. In particular, [35] is an up to date review that considers many validation metrics in the cases of data and model certainty, data uncertainty and model certainty, and data and model uncertainty. Further, they comment on the importance of model validation while admitting its contextual difficulties because often validation is “strongly application, quantity of interest, and modeler, dependent”. Our article addresses some of these contextual difficulties by taking a bottom up approach.

To assist the comparison of the positive and negative aspects of (most) the above validation metrics, reference [34] outlines six “desirable validation criteria” that a validation metric might have (they extend [27, 15]). One conclusion from [34] is that none of the available metrics simultaneously satisfy all six desirable validation criteria. We summarize the most important features of the desirable validation criteria with the following validation criterion:

1. A validation metric should be a statistically quantified quantitative measure (as opposed to a qualitative measure) of the *agreement* between (general) model predictions and data pairs, in the presence or absence of uncertainty.

The desire for objectivity, “that a metric will produce the same assessment for every analyst independent of their individual preferences” [34], is difficult to satisfy because there are no rules in place to guide a modeler toward selecting one validation metric over another. For this reason the individual might simply choose a metric based on their preferences, or worse, be tempted to base their decision on which validation metric gives them the most favorable evaluation. Given individuals may choose different validation metrics for the same model-data pair, it is possible for individuals to impose accuracy requirements that are incompatible with one another and arrive at different conclusions regarding the validity of the same model-data pair. As the final goal is objectivity, when possible, a map between the accuracy requirements should be constructed such that the validation metrics yield consistent evaluations of model-data validity when applicable.

Further, Liu et al. [34] suggest that there is no agreed upon unified model-data comparison function. Even including the results of this article, we expect this statement to hold as it is extremely difficult to guess the prior information about the utility of a model an analyst may be required to include in the validation of a model-data pair. For instance given arbitrary data “What features of the data are relevant to capture with a model?”, “Of these features, are some more relevant than others?”, and “What accuracy is required for the model to be valid?”. *Agreement* and *validation* are ultimately human-made concepts designed for the purpose of expressing that “in general, not every feature or statistic between a model-data pair need to be equal to conserve the utility of the model”. For some model-data pairs, all that may be required is that the model and data averages closely match within uncertainty, while for others, one may require that the model can accurately reproduce the probability distribution of a data set as a whole (as one would do to physically model a noisy measurement device). Given the wide variety of data and the large number of different inferences (and thus models and hypotheses) that one may be interested in drawing from a given data source, i.e. the context of the model-data pair, we do not expect any single set of comparison functions, statistics, or values to be equally relevant and maximally useful for all possible model-data contexts. This, however, does not stop us from quantifying the validity of a model-data pair given any arbitrary comparison function and with any arbitrary definition of agreement.

In this article, we construct and propose the “Bayesian Validation Metric” (BVM) as a general model validation and testing tool. We design the BVM to adhere to the desired validation criterion (1.) by using “four BVM inputs”: the model and data comparison values, the model output and data pdfs, the comparison value function, and the agreement function. The comparison value function is a function of model output and data comparison values that provides the desired quantitative comparison measure, e.g. square difference. Using the model output pdf and the data pdf, the value of the comparison value function is statistically quantified. In turn, the agreement function provides an accept/reject rule and effectively wraps the previous three BVM inputs together to give the BVM. From this, the BVM outputs “the probability the model and data agree”, where *agreement* is a user defined Boolean function that meets, or does not meet, accuracy

requirements between model and data comparison values. Thus, the BVM meets the desired validation criterion (1.) for arbitrary comparison value functions, arbitrary definitions of agreement, and in principle for arbitrary data types such as integers, vectors, tensors, strings, pictures, or others.

The BVM can be used to represent all of the aforementioned validation metrics as special cases. This allows us to compare and contrast the validation metrics from the literature in a new light. We find the conditions under which several of the current validation metrics are effectively equal to one another, which improves the objectivity of the current validation procedure. In brief we find that the frequentist metric (using natural definitions of agreement) is equal to the reliability metric and the probabilities from Bayesian model testing are equal to the probabilities of the improved model reliability metric [26] when one demands exact equality of the (uncertain) model-data comparison values. Because probability can represent both certain and uncertain situations, so can the BVM. Thus, these “special case” metrics can be generalized to quantify certain or uncertain cases, and even be combined into more complex validation requirements using the BVM framework. Thus, the BVM provides a standardized framework to improve, generalize, or further quantify these validation metrics.




Figure 1: This is the depiction of a common model validation scenario. The model line is trained on noisy data (not depicted in the figure) and is to be compared to a set of validation data. As both the model line and the data are uncertain in general, any quantitative measure (i.e. the comparison function) between these comparison values inherits this uncertainty. Thus, any accept/reject rule on the basis of these uncertain comparison function values is uncertain as well. A visual inspection of this graph seems to indicate, up to statistical fluctuation, that the comparison values of the model and data more or less (or probably) agree, but this intuitive measure has yet to be quantified. Graphic adapted from [38].

By constructing the “BVM ratio”, we generalize the Bayesian model testing framework [29]. In the Bayesian model testing framework, one constructs the Bayes ratio to rank models according to the ratio of their posterior probabilities given the data. We show that these posterior probabilities are equal to a special case of the BVM under the definition of agreement that requires these uncertain model outputs and data

to match exactly. Thus, nothing prevents us from extending the logic used in the Bayesian model testing framework to our framework and we construct the BVM ratio for the purpose of model selection under arbitrary definitions of agreement, i.e. for arbitrary validation scenarios.

The remainder of the article continues as follows. In Section 2 we construct the BVM by following our validation criterion. Through some edge cases, we show that the BVM satisfies both the six desirable validation criteria from [34] as well as our validation criterion (1.) in Section 3. In Section 4, we summarize the results derived in Appendix A. There, we incorporate all of the above standard validation metrics as special cases of the BVM, draw relationships between several of the validation metrics, provide improvements and generalizations to these metrics as is suggested by the functional form of the BVM, and construct the BVM ratio. In Section 5 we represent three novel validation metrics using the BVM and compare them to similar metrics from the literature. Due to the proven capacity of our framework, we consider the BVM to be a general model validation and testing tool.

2 The Bayesian Validation Metric

For the remainder of the article we will use the following notation and language. We will let \hat{y} denote the output of a model, \hat{Y} a set of model outputs, y a data point, and Y a set of data points. The proposition M essentially stands for “the model” or “coming from the model”, and D stands for “the experiment” or “coming from the experiment”. We let \hat{z} and z represent the comparison quantities of interest, which pertain to the model and the data respectively. Further, we let the comparison quantities take general forms, such as multidimensional vectors, functions, or functionals (e.g. output values, expectation values, pdfs, ...), so we can represent any such pair of quantities we may wish to compare between the model and the data. When we refer to “the four BVM inputs” we mean: the comparison values (\hat{z}, z) , the model output and data pdf $\rho(\hat{z}, z|M, D)$, the comparison value function $f(\hat{z}, z)$, and the agreement function $B = B(f(\hat{z}, z))$. The (denoted) integrals may be integrals or sums depending on the nature of the variable being summed or integrated over, which is to be understood from the discrete or continuous context of the inference at hand. The dot “ \cdot ” represents standard multiplication, which is mainly used to improve aesthetics.

Performing uncertainty propagation through a model results in a model output probability (density) distribution $\rho(\hat{y}|M, D)$ that ultimately we would like to validate by comparing it to an uncertain validation data source $\rho(y|D)$, to see if they agree (as depicted in Figure 1). The immediate question is, however, “What values do we want to *compare* and what do we mean by *agree*?”. Given the wide variety of data and the large number of different inferences (and thus models and hypotheses) that one may be interested in drawing from a given data source, i.e. the context of the model-data pair, we do not expect any single set of comparison functions to be equally relevant and maximally useful for all possible model-data contexts. In light of this, we instead quantify the validity of a model-data pair given any arbitrary comparison value function and with any arbitrary definition of agreement.

2.1 Derivation

Here we will begin constructing the Bayesian Validation Metric (BVM). To capture the concept of what we might mean by *agree*, we define \hat{z} and z to agree, A , when the Boolean expression, B , is true. Both A and B are defined by the modeler and their prior knowledge of the context of the model-data pair. Naturally then, the agreement function $B = B(f(\hat{z}, z)) = B(\hat{z}, z)$ is some function or functional of a comparison value function $f(\hat{z}, z)$.

Given the values of \hat{z} and z are known, i.e. certain, we quantify *agreement* using a probability distribution that assigns certainty,

$$p(A|\hat{z}, M, z, D) = \Theta(B(\hat{z}, z)). \quad (1)$$

The indicator function $\Theta(B)$ is defined to equal unity if B evaluates to “true” (i.e. “agreeing”) and equal to zero otherwise. Thus in the completely certain case, we are certain as to whether the model and data

comparison values ‘agree, or do not agree, as defined by B and the deterministic evaluation of $f(\hat{z}, z)$.¹ We will call $p(A|\hat{z}, M, z, D)$ the “agreement kernel”.

Given that in general the comparison values are uncertain, and quantified by $\rho(\hat{z}, z|M, D)$, the probability the comparison values agree, as defined by B and $f(\hat{z}, z)$, is equal to,

$$p(A|M, D) = \int_{\hat{z}, z} p(A|\hat{z}, M, z, D) \cdot \rho(\hat{z}, z|M, D) d\hat{z} dz, \quad (2)$$

$$\xrightarrow{\text{ind.}} \int_{\hat{z}, z} \rho(\hat{z}|M, D) \cdot \Theta(B(\hat{z}, z)) \cdot \rho(z|D) d\hat{z} dz, \quad (3)$$

which is a marginalization over the spaces of (\hat{z}, z) .² Equation (2) is the general form of the Bayesian Validation Metric (BVM). Because A is discrete, the BVM is a probability rather than a probability density and it therefore falls in the range $0 \leq p(A|M, D) \leq 1$. Equation (3) explicitly assumes that the uncertainty in the data is independent of the model, i.e. $\rho(z|M, \hat{z}, D) = \rho(z|D)$, that the data D does not take \hat{z} or the model M (that it is currently being compared to) as inputs.³ This is a relatively common scenario so it is stated explicitly. The BVM may be given a geometric interpretation as the projection of two probability vectors (potentially of unequal length) in a space whose overlap is defined by the agreement kernel – this is an inner product if $B(\hat{z}, z)$ is symmetric in its arguments. The BVM may be computed using any of the well known computational integration methods.

2.2 An identical representation

In some cases, it is useful to work directly with the probability density $\rho(f|M, D)$, which quantifies the probability the comparison value function $f(\hat{z}, z)$ takes the value f due to uncertainty in its inputs. This pdf is independent of any user defined accuracy requirement. We will call this pdf the comparison value probability density, which is equal to,

$$\rho(f|M, D) = \int_{\hat{z}, z} \delta(f - f(\hat{z}, z)) \cdot \rho(\hat{z}, z|M, D) d\hat{z} dz. \quad (4)$$

This is the net uncertainty propagated through the comparison value function $f(\hat{z}, z)$ from the uncertain model and data comparison values. All of the expectation values that are associated with f may be generated from this pdf.

If one imposes an accuracy requirement with a Boolean expression $B = B(f)$ (i.e. defining agreement according to the value of f), the resulting accumulated probability is the BVM. That is, the BVM, i.e. equation (2), may equally be expressed as,

$$p(A|M, D) = \int_f \rho(f|M, D) \cdot \Theta(B(f)) df, \quad (5)$$

which is proven through substitution and marginalization over f ,

$$\begin{aligned} p(A|M, D) &= \int_f \left(\int_{\hat{z}, z} \delta(f - f(\hat{z}, z)) \cdot \rho(\hat{z}, z|M, D) d\hat{z} dz \right) \cdot \Theta(B(f)) df \\ &= \int_{\hat{z}, z} \Theta(B(\hat{z}, z)) \cdot \rho(\hat{z}, z|M, D) d\hat{z} dz. \end{aligned} \quad (6)$$

¹This binary yet probabilistic definition of agreement turns out to be completely satisfactory for our current purposes. As is briefly discussed later, the sharp boundaries of the indicator function can be smoothed out without employing fuzzy logic by allowing parameters in the Boolean function to themselves be uncertain and marginalized over.

² Recall that the propositions in the probability distributions $\rho(z|D)$ and $\rho(\hat{z}|M, D)$ are completely arbitrary (in some cases requiring propagation from $\rho(y|D)$ and $\rho(\hat{y}|M, D)$), they could be both continuous, discrete (with order), categorical variables (no well defined order, e.g. strings, pictures,...), or a mix.

³In a controls system this may not be the case, as the model may interact with the system of interest. In such a case this constraint may be lifted and one should use (2) instead. The joint probability $\rho(\hat{z}, z|M, D)$ can be used to account for the correlations between the model (the controller or reference) and the data (the measured response of the system being controlled) in a controls setting in principle.

2.3 Importance

The BVM allows the user to, in principle, quantify the probability the model and the data agree with one another under arbitrary comparison value functions and with arbitrary definitions of agreement. The BVM can therefore be used to fully quantify the probability of agreement between arbitrary model and data types using novel or existing comparison value functions and definitions of agreement. Thus, the problem of model-data validation may be reduced to the problem of finding the four BVM inputs in any model validation scenario.

2.4 Statistical Responsibility

When using the BVM framework, one should practice statistical responsibility by explicitly stating the definition of agreement that is implemented in the validation procedure. Although the flexibility of the BVM framework is a feature, different validation metrics often have different amounts of tolerance as what constitutes “agreement”. Agreement according to one metric does *not* in general imply agreement according to another. Overly tolerant definitions of agreement have little resolution power and can only be used responsibly if a large degree of non-exactness between the model and data is permissible. In principle, the definition of agreement should be just as strict or stricter than it needs to be. By explicitly stating the definition of agreement along side the BVM value, $p(A|M, D) \equiv p(A|M, D, B)$, one avoids statistical misrepresentation by not hiding one’s definition of agreement.

3 Meeting the desirable validation criterion

First we will describe how the BVM, equations (2) and (5), precisely match our validation criterion (1.). As can be seen by equation (4), incorporated into the BVM is a statistically quantified quantitative measure that compares data and model outputs, $\rho(f|M, D)$. However, this pdf is in some sense lacking a context pertaining to the model-data pair. Not until an accept/reject rule is imparted on $\rho(f|M, D)$ does one define what is meant by *agreement* in the model-data context. Thus, the BVM only becomes the probability of *agreement* between the data and the model when the agreement function is also incorporated. The four BVM inputs are therefore adequate to satisfy (1.) as the BVM is a “statistically quantified quantitative measure (f) of agreement $p(A|M, D)$ between model predictions and data pairs (\hat{z}, z) , in the presence or absence of uncertainty $\rho(\hat{z}, z|M, D)$ ”.

There are a few more BVM concepts worth discussing before moving forward. We will show that the BVM is capable of handling general multidimensional model-data comparisons and that there are no conceptual issues when agreement is exact, i.e. B is true iff $\hat{z} = z$, in the certain and uncertain cases. We will then make comments on the sense in which the BVM adheres to the full set of six desirable validation criteria given in [34] by discussing the criteria that are underrepresented in (1.).

3.1 Compound Booleans

Because Boolean operations between Boolean functions results in a Boolean function itself, the BVM is capable of handling multidimensional model-data comparisons. We will call a Boolean function with this property a “compound Boolean”. A compound Boolean function results from *and*, \wedge , conjunctions and *or*, \vee , disjunctions between a set of Boolean functions, e.g.,

$$B(\{B_i\}) = \left(B_1(\hat{z}, z) \vee B_2(\hat{z}, z) \right) \wedge \left(B_3(\hat{z}, z) \vee B_4(\hat{z}, z) \right) \dots, \quad (7)$$

where each $B_i(\hat{z}, z) = B_i(f_i(\hat{z}, z))$ may use a different comparison function $f_i(\hat{z}, z)$. Compound Booleans using conjunctions quantifying the validity of entire model functions (random fields and/or multidimensional vectors) by assessing agreement between each of the model-data comparison field points simultaneously, i.e. over the comparison points 1 *and* points 2 *and* so on. The compound Booleans may be factored into their constituting Boolean functions using the standard product and sum rules of probability theory after being

mapped to probabilities with the agreement kernel. One should be weary when defining an *and* Boolean that, if one of the Booleans is false, then the entire Boolean is false. If this strict “all or nothing” validation requirement is not needed then other more flexible definitions of agreement may be instantiated instead (see “BVM Examples”).

3.2 The BVM under the conditions of exact agreement

We can calculate the BVM under the conditions of exact agreement in the completely certain and uncertain cases. Because the BVM is a probability rather than a probability density, the agreement kernel falls in the range $[0, 1]$. Under the conditions of exact agreement, that B is only true when $\hat{z} = z$, the agreement kernel is $\Theta(B) = \Theta(\hat{z} = z) = \delta_{\hat{z}, z}$, which is the Kronecker delta, i.e. it is 0 or 1, but has continuous labels. As it is uncommon to deal with Kronecker delta’s having continuous labels under integration, we will show that the BVM gives reasonable results under the condition of exact agreement in the complete certainty as well as in the general uncertain case.

Complete certainty and exact agreement Complete certainty is represented using Dirac delta pdf functions over the model and data comparison values. This gives the BVM,

$$\begin{aligned} p(A|M, D) &= \int_{\hat{z}, z} \rho(\hat{z}|M, D) \Theta(\hat{z} = z) \rho(z|D) d\hat{z} dz \\ &= \int_{\hat{z}, z} \delta(\hat{z} - z') \cdot \delta_{\hat{z}, z} \cdot \delta(z - z') d\hat{z} dz, \end{aligned} \quad (8)$$

where we are considering the model-data pair to agree iff the comparison values are exactly equal. Using the sifting property of the Dirac delta function, we find the reasonable result that,

$$p(A|M, D) = \delta_{\hat{z}', z'}, \quad (9)$$

which is equal to unity iff \hat{z}' and z' , the definite values of \hat{z} and z , are equal.

Uncertainty and exact agreement In the uncertain case under the condition of exact agreement, the BVM is

$$p(A|M, D) = \int_{\hat{z}, z} \rho(\hat{z}|M, D) \cdot \delta_{\hat{z}, z} \cdot \rho(z|D) d\hat{z} dz. \quad (10)$$

We will do the following trick to correctly interpret this integral. We will first let $B(\epsilon)$ be true if $z - \epsilon \leq \hat{z} \leq z + \epsilon$ and then take the limit as $\epsilon \rightarrow 0^+$ such that $\lim_{\epsilon \rightarrow 0^+} B(\epsilon) \rightarrow B$ when appropriate. With this Boolean expression, the BVM is,

$$\begin{aligned} p(A|M, D, \epsilon) &= \int_{\hat{z}, z} \rho(\hat{z}|M, D) \Theta(z - \epsilon \leq \hat{z} \leq z + \epsilon) \rho(z|D) d\hat{z} dz \\ &= \int_z \rho(z|D) \left(\int_{z-\epsilon}^{z+\epsilon} \rho(\hat{z}|M, D) d\hat{z} \right) dz. \end{aligned} \quad (11)$$

In the limit $\epsilon \rightarrow 0^+$, the term in the parenthesis $\int_{z-\epsilon}^{z+\epsilon} \rho(\hat{z}|M, D) d\hat{z} \rightarrow p(\hat{z} = z|M, D) = \rho(\hat{z} = z|M, D) d\hat{z}$ by the definition of probabilities. This gives,

$$\begin{aligned} p(A|M, D) &= \int_z \rho(z|D) \left(p(\hat{z} = z|M, D) \right) dz = \left(\int_z \rho(\hat{z} = z|M, D) \rho(z|D) dz \right) d\hat{z} \\ &\equiv \rho(\hat{z} \equiv z|M, D) d\hat{z} = p(\hat{z} \equiv z|M, D), \end{aligned} \quad (12)$$

which is understood to be the sum of the model and the data probabilities that jointly output exactly the same values. We see that the BVM in this case is proportional to $d\hat{z}$, $p(A|M, D) \rightarrow \rho(\hat{z} \equiv z|M, D) d\hat{z}$, in the general

case of exact agreement, and therefore the BVM goes to zero unless the pdf $\rho(\hat{z} \equiv z|M, D) \propto \delta(\dots)$. Thus, we recover the standard logical result for probability densities $p(x) = \rho(x)dx \rightarrow 0$ unless it is offset by $\rho(x) \propto \delta(x)$. This result is easily generalized to the dependent case using $\rho(\hat{z}, z|M, D) = \rho(z|M, D)\rho(\hat{z}|z, M, D)$. The result, equation (12), is no more surprising than (9) in principle.

Due to the vast number of possibilities for continuous valued variables, having a pathological definition of exact agreement between continuous variables does not occur in practice. In a computational setting, $dx \rightarrow \Delta x$ becomes a finite difference and these infinitely improbable agreement conceptual issues are avoided. The Bayesian model testing framework avoids these issues by evaluating posterior odds ratios, in which case the measures, $d\hat{z}$, drop out.

3.3 Meeting underrepresented validation criteria

In this subsection we will discuss how the BVM also meets the validation criteria found in [34]. This is done by using the derived general and special cases of the BVM for each of the criteria which are underrepresented in (1).

Perhaps the primary underrepresented criterion from [34] is their second. It states that “..the criteria used for determining whether a model is acceptable or not should not be a part of the metric which is expected to provide a quantitative measurement only.” We argue that the functional form of the BVM presented in equation (5) clearly demonstrates this feature as it factors into $\rho(f|M, D)$ and $\Theta(B(f))$. The comparison function $f(\hat{z}, z)$ represents the “objective quantitative measure” from their first criterion that is separate from the accept/reject rule, which is our agreement function $B(f)$ – both of which require definition to ultimately evaluate the validity of a model. We see it as advantageous to quantify the probability the model is accepted or rejected through $B(f)$ due to the uncertainty in the value of f , which is the general case, and which gives the BVM as the result. As all of the validation metrics presented in [34] (and more) will be shown to be representable with the BVM, and thus placed on the same footing, we find our language of “comparison function” and “agreement function” to ultimately be more useful than a language that only considers comparison functions (without accept/reject rules) to be the validation metrics.

The third criteria in [34] is that ideally the metric should “degenerate to the value from a deterministic comparison between scalar values when uncertainty is absent”. This is indeed the case as can be seen in equation (1) or in equations (2) and (5) by utilizing Dirac delta pdfs similar to their application in equation (8).

The fifth desirable validation criteria in [34] states that artificially widening probability distributions should not lead to higher rates of validation. They find all but the frequentist metric to have this undesired feature; however, we later see that the frequentist metric may be considered a special case of the reliability metric (when reasonable accuracy requirements are imposed), meaning artificial widening can lead to higher rates of validation for more general instances of the frequentist metric.

Further, we argue that artificially introducing uncertainty for the express purpose of passing a validation test is indistinguishable from scientific misconduct. If there is objective reason to include more uncertainty into the analysis or if the circumstance for what constitutes validation has changed due to a change of context – and it happens to improve the rate of validation – so be it. This is a different context, model, or state of uncertainty than was originally proposed so different rates of acceptance should be expected. Reducing the uncertainty of either the data or the model (the inputs) through additional measurements or changing the model may later prove the model valid or invalid when it may have been initially accepted. Thus, to meet this validation criteria, we simply assume the user is not engaging in scientific misconduct.

Finally, due to the results of the “Compound Booleans” section, their sixth criterion is met. Because the BVM (2) can be used to assess single or multidimensional controllable settings (see footnote number 3.) we can perform global function validity (in or out of a controls setting). As they note, “This last feature is critical from the viewpoint of engineering design”.

Thus, the BVM satisfies both our validation criterion and the six desirable validation criteria outlined in [34]. This was accomplished by representing model-data validation as an inference problem using the four BVM inputs.

4 Representing and generalizing the known validation metrics with the BVM

This section is a review of the material found in Appendix A. The following validation metrics will be represented with the BVM, which are then improved, generalized, and/or commented on: reliability/probability of agreement, improved reliability metric, frequentist, area metric, pdf comparison metrics, statistical hypothesis testing, and Bayesian model testing.

4.1 Representing the known validation metrics with the BVM

Table 1 outlines the values of the four BVM inputs that result in the BVM representing each of the well known validation metrics as special cases. The following notation is used for the comparison values (\hat{z}, z) . The brackets $\langle \dots \rangle$ denote expectation values, μ 's denote averaged values, \hat{y}, y denote single values, \hat{Y}, Y denote multidimensional (or many valued) values, F_y, F_y denote cumulative distribution functions ($F_y = \int_{-\infty}^{\hat{y}} \rho(\hat{y}|M, D) dy$), $S_{\hat{y}}, S_y$ denote test statistics, and $[-c_\alpha, c_\alpha]$ denotes the $1 - \alpha$ confidence interval of the data. In the agreement function column, an element listed as $B(f)$ means the creators of the metric intentionally left the definition of agreement unspecified; however, it is natural to assume it is a function of the comparison function f .

Table 1 shows the specification of the four BVM inputs that give the other validation metrics as special cases. It also summarizes some of the similarities and difference between the known validation metrics. In particular, by looking at the validation metrics with the same type of comparison values, i.e. the reliability and frequentist or the improved reliability and Bayesian model testing, we can compare them directly. We see that if one lets the frequentist metric allow for more general input probability distributions and the use of a reasonable agreement function (i.e., $B(f)$ is true if $|f| < \epsilon$), then the frequentist metric is the reliability metric. Further, in Bayesian model testing, if the agreement function $B(f)$ is loosened to accept $f < \epsilon$, than the pdf's that appear in the Bayesian model testing framework are equal to the improved reliability metric. This information improves the objectivity of the current validation procedure because we now have a map between validation metrics that were originally thought to be different.

Table 1: Specification of the four BVM inputs that give the other validation metrics as special cases.

	Comp. Values		Probs.		Comp. Func.	Agree. Func.
BVM	\hat{z}	z	$\rho(\hat{z} M, D)$	$\rho(z D)$	$f(\hat{z}, z)$	$B(f)$
Reliability	$\langle \hat{y} \rangle$	μ_y	$\rho(\langle \hat{y} \rangle M, D)$	$\rho(\mu_y D)$	$ \langle \hat{y} \rangle - \mu_y $	$f < \epsilon$
Imp. Reli.	\hat{Y}	Y	$\rho(\hat{Y} M, D)$	$\rho(Y D)$	$ \hat{Y} - Y $	$f < \epsilon$
Frequentist	$\langle \hat{y} \rangle$	μ_y	$\delta(\langle \hat{y} \rangle - \langle \hat{y} \rangle')$	Stud. t	$\langle \hat{y} \rangle - \mu_y$	$B(f)$
Area	F_y	F_y	$\delta(F_y - F_y')$	$\delta(F_y - F_y')$	$\int_y F_y - F_{y=\hat{y}} dy$	$B(f)$
Pdf Comp.	ρ_M	ρ_D	$\delta(\rho_M - \rho_M')$	$\delta(\rho_D - \rho_D')$	$G(\rho_D \rho_M)$	$B(f)$
Stat. Hyp.	$S_{\hat{y}}$	S_y	$\rho(S_{\hat{y}} M = D)$	$\rho(S_y D)$	$S_{\hat{y}}$	$S_{\hat{y}} \in [-c_\alpha, c_\alpha]$
Bayes Model	\hat{Y}	Y	$\rho(\hat{Y} M, D)$	$\rho(Y D)$	$ \hat{Y} - Y $	$f = 0$

The column headings are the four BVM input values: Comparison Values (\hat{z}, z) , Probabilities $(\rho(\hat{z}|M, D), \rho(z|D))$, Comparison Function $f = f(\hat{z}, z)$, and the Boolean Agreement Function $B(f)$. The row headings read: Reliability, Improved Reliability, Frequentist, Area, Pdf Comparison Metrics, Statistical Hypothesis Test, and Bayesian Model Test. The denoted data probability for the average in the frequentist metric, Stud. t, is the Student t distribution.

Table 2: BVM representation of the special cases using the f 's specified in Table 1.

	BVM
BVM	$\int_f \rho(f M, D) \cdot \Theta(B(f)) df$
Reliability	$\int_f \rho(f M, D) \cdot \Theta(f < \epsilon) df = r$
Imp. Reli.	$\int_f \rho(f M, D) \cdot \Theta(f < \epsilon) df = r_i$
Frequentist	$\int_{\mu_y} \rho(\mu_y D) \cdot \Theta(B(f(\langle \hat{y} \rangle', \mu_y))) d\mu_y$
Area	$\Theta(B(f(F'_y, F'_y)))$
Pdf Comp.	$\Theta(B(f(\rho'_M, \rho'_D)))$
Stat. Hyp.	$\int_{S_{\hat{y}}} \rho(S_{\hat{y}} M = D) \cdot \Theta(S_{\hat{y}} \in [-c_\alpha, c_\alpha]) dS_{\hat{y}} = 1 - \alpha$
Bayes Model	$\int_f \rho(f M, D) \cdot \Theta(f = 0) df = p(\hat{Y} \equiv Y M, D)$

Table 2 shows the resulting BVM using the specifications listed in Table 1. The value r is the standard notation for the reliability metric [24] and we use r_i for the improved reliability metric [26]. The BVM represents each of the known validation metrics as a probability of agreement between the model and the data from equation (2). As no agreement function is specified directly for the frequentist and area metric, the problem is under constrained so the agreement functions are left as general functions over the comparison function $B(f)$. Thus, for any chosen agreement function, the BVM quantifies their probability of agreement. The remaining metrics all do specify (or indicate) an agreement function, and thus, have specified all of the information required to compute the BVM.

The statistical hypothesis test is perhaps a bit out of place among the validation metrics. First, note that the comparison function for statistical hypothesis testing is not a function of both the data and the model. Further note, the model pdf used for statistical hypothesis testing assumes the null hypothesis is true, which in our language is the assumption that $\rho(S_y|M, D) = \rho(S_y|M = D)$, i.e. that the pdf of the model is equal to the pdf of the data. This shows how statistical hypothesis is a bit out of place here among the validation metrics because here we are attempting to validate a model, usually with its own quantified pdf, rather than, perhaps irresponsibly, assuming it is equal the data pdf before validating that to be the case. This causes standard statistical hypothesis pitfalls, such as type I (rejecting the null hypothesis when it is true) and type II errors (failing to reject the null hypothesis when it is false), to be carried over into BVM, which is unwanted. Several comments are made in Appendix A.5 on this issue.

A perhaps surprising result is the proposed functional form of the BVM that represents Bayesian model testing $p(A|M, D) = p(\hat{Y} \equiv Y|M, D)$, which is the Bayesian evidence. This is the probability the uncertain model and data happen to output exactly the same values. Usually what is discussed when reviewing Bayesian model testing is the Bayes posterior odds ratio, i.e. the ‘‘Bayes Ratio’’,

$$R = \frac{p(M|Y)}{p(M'|Y)} \propto \frac{p(Y|M)}{p(Y|M')},$$

which tests one model M , i.e. for validation, against another model M' . However, in validation metric problems, we are first interested in considering the validation of a single model – the ratio is an extra bit of inference. In Appendix A.6 we show that the BVM result of $p(\hat{Y} \equiv Y|M, D)$ is exactly what we mean by $p(Y|M)$ in the numerator of the Bayes factor,⁴ which effectively quantifies the validation of a single model against data Y , all quantified under uncertainty.

4.2 Generalizations and Improvements to the known validation metrics with the BVM

The BVM offers several avenues to either generalize or improve many of the metrics. The types of generalizations the BVM offer pertain to generalizing the comparison values, comparison functions, definitions of

⁴It should be noted that our notation for D differs from the notation typically used in Bayesian model testing. Their D is equal to our data Y , while our D refers to context ‘‘as having come from the data or experiment rather than the model’’.

agreement, and/or generalizing deterministic comparison values and metrics to the uncertain case. These generalizations are only useful if quantitative statements can be made on their behalf – in such a case, these generalizations are improvements. We will give a brief review of the improvements we found below, but the full discussion is located in Appendix A. By making generalizations or improvements to each of the known validation metrics as implied by the BVM, each metric can be made to satisfy our validation criterion as well as the six desirable validation criteria in [34], due to the results of Section 3.

Appendix A.1 uses the BVM to show that the reliability metric and the improved reliability metric can be generalized to compare values without a unique order, such as strings, in principle. This involves creating an agreement function over sets of values (such as synonymous sets of strings), rather than in a continuous interval, that may be considered to “agree”.

Appendix A.2 derives the frequentist validation metric and generalizes it to the case where both the model and data expectation values are uncertain. The frequentist metric assumes the model outputs are known with certainty, which may or may not be true. If a model is stochastic, the model pdfs may be estimated with Monte Carlo or other uncertainty propagation methods that quantify the pdf directly.

Appendix A.3 shows that the area metric may be cast as a special case of the BVM. The area metric involves quantifying the difference between model and data cumulative distributions on a point to point basis; thus, the comparison values (\hat{z}, z) are cumulative distributions themselves. The comparison values are assumed to be known with complete certainty, which in the case of cumulative distributions of data is often difficult to argue. Any quantifiable uncertainty in the cumulative distributions may be integrated over, which generalizes the area metric to situations when the model and/or the data cumulative distributions are uncertain. A drawback is that the BVM in these cases may be very computationally intensive and would likely need to be approximated using a random sampling or discretization scheme. A binned pdf metric is put forward to potentially reduce the computational complexity toward quantifying this generalized area validation metric. This applies similarly to the pdf comparison metrics in Appendix A.4.

In Appendix A.5 we invent an improved statistical hypothesis test using the BVM that takes into account both model and data pdfs called the “statistical power BVM”. Because in principle we have a model output pdf in model validation problems $\rho(\hat{y}|M, D)$, we can use it (in place of assuming the null hypothesis is true) to remove the possibility of both type I and type II errors.

In the statistical power BVM, the model and the data are defined to agree if their test statistics both lie within one another’s confidence intervals (or “confidence sets” as explained in Appendix A.5). The statistical power BVM becomes the product of the statistical powers of the model and data, denoted $p(A|M, D) = (1 - \beta_D(\hat{\alpha})) \cdot (1 - \beta_M(\alpha))$ in equation (46). Further comments are made about how systematic error (defined as when a test statistic lies outside of its *own* confidence interval) may be removed.

It is concluded that the statistical power BVM has a relatively low resolving power compared to other BVMs. This is because large confidence intervals imply large tolerance intervals for acceptance. For this reason, statistical hypothesis testing should only be used for validation in situations when a high degree of nonexactness between model and data test statistics is permissible and the pdf’s have very thin tails. This BVM does, however, have a greater resolution than the classical hypothesis test as was proved in the appendix and will be demonstrated next section.

Appendix A.6 finds that Bayesian model testing has the highest possible resolving power because the model and the data are defined to agree only if their values are exactly equal. This is the reverse of what was concluded about statistical hypothesis testing.

Further in Appendix A.6, we argue that, analogous to the Bayesian model testing framework, nothing prevents us from constructing what we call the BVM factor. The BVM factor is,

$$K(B) = \frac{p(A|M, D, B)}{p(A|M', D, B)}, \quad (13)$$

which is a ratio of the BVMs of two models under arbitrary definitions of agreement B . Using Bayes Theorem, $p(M|A, D, B) = \frac{p(A|M, D, B)P(M|D, B)}{P(A|D, B)}$, we may further construct the BVM ratio,

$$R(B) = \frac{p(M|A, D, B)}{p(M'|A, D, B)} = \frac{p(A|M, D, B)P(M|D, B)}{p(A|M', D, B)P(M'|D, B)} = K(B) \frac{P(M|D, B)}{P(M'|D, B)}, \quad (14)$$

for the purpose of comparative model selection under a general definition of agreement B . The ratio $\frac{P(M|D,B)}{P(M'|D,B)}$ is the ratio prior probabilities of M and M' , which analogous to Bayesian model testing, if there is no reason to suspect that one model is a priori more probable than another, one may let $\frac{P(M|D,B)}{P(M'|D,B)} = 1$, and then $R(B) \rightarrow K(B)$ in value.

Thus, using the BVM ratio, we can perform *general model validation testing under arbitrary definitions of agreement and with any reasonable set of comparison functions*. The BVM ratio therefore generalizes the Bayesian model testing framework. This will be utilized in next section.

Finally we wanted to add a note about how one may mitigate the sharpness of the indicator function without using fuzzy logic while also allowing close models to be somewhat accepted. As we have seen, it is natural to use a threshold Boolean parameter ϵ to help define the boundary of agreement through $B(f \leq \epsilon)$. Such a BVM takes the form,

$$p(A|M, D, \epsilon) = \int_f \rho(f|M, D) \cdot \Theta(B(f \leq \epsilon)) df, \quad (15)$$

where $\Theta(\dots)$ instantaneously drops to zero for $f > \epsilon$. One may soften the boundary by allowing ϵ itself be an uncertain quantity, which means one allows their definition of agreement to be somewhat uncertain (which often times can be reasonably claimed). As an example, let this uncertainty be $\rho(\epsilon) = \lambda \exp(-\lambda(\epsilon - \epsilon'))$ for $\epsilon' > \epsilon$ and zero otherwise, where λ is positive. Marginalizing over this BVM then gives,

$$\begin{aligned} p(A|M, D) &= \int_{\epsilon} p(A|M, D, \epsilon) \rho(\epsilon) d\epsilon \\ &= \int_f \rho(f|M, D) \cdot \left[\Theta(B(f \leq \epsilon')) + \Theta(B(f > \epsilon')) e^{-\lambda(f - \epsilon')} \right] df, \end{aligned} \quad (16)$$

which allows some f 's to be accepted outside the agreement region defined by $f \leq \epsilon'$, but with a exponentially decaying probability. Other potentially useful ϵ pdfs include, but are not limited to: negative slope linear, Gaussian, or decaying sigmoidal. None of these ϵ type distributions were needed to obtain the results of the previous sections explicitly; however, these types of assumptions may have been part of the decision process made implicitly by a practitioner while performing model validation.

5 BVM Examples

In this section we invent and quantify three novel validation metrics using the BVM to highlight the conceptual clarity, flexibility, and capacity of our framework.

5.1 The Statistical Power BVM

Here we consider the statistical power BVM proposed in A.5 and reviewed in the previous section. This metric defines agreement as occurring when both the model and data comparison values are within one another's confidence intervals, simultaneously. The BVM for this metric is the product of the statistical powers of the model and the data $p(A|M, D) = (1 - \beta_D(\hat{\alpha})) \cdot (1 - \beta_M(\alpha))$, which is calculated in equation (46). We contrast this with the standard statistical hypothesis test that, after assuming the model is correct, $M = D$, finds the probability the model lies within the data's confidence interval equal to $1 - \alpha$. In the statistical hypothesis testing, one then proceeds to check the actual model output and speculates about type I and II errors. As discussed in the Appendix, we do not assume $M = D$ before validation and therefore type I and II errors are avoided. Rather, we let the statistical power BVM decide whether or not the model is valid. This provides a more informative validation procedure.

Figure 2 depicts a typical statistical hypothesis test scenario that is designed to check the validity of an uncertain model average prediction $\hat{\mu}$ (in blue) against an uncertain data average prediction μ (in red). The data's μ is t -distributed the same in each subfigure according to $T(\bar{y}, n - 1, \bar{s}) = T(0, 10, 1.75)$ where (\bar{y}, n, \bar{s})

are the sample mean, number of collected data points, and the sample standard deviation, respectively. Each row depicts a normal distribution model centered at 0, but with increasing model variance per row.




Figure 2: The shaded regions in Column A depict the 95% confidence interval of each distribution, respectively. Because the data distribution is the same in each figure and because the statistical hypothesis test is independent of the proposed model due to assuming the hypothesis $M = D$, each model is equally valid by that test when it is clear that the model in row B is preferable. The shaded regions in Column B depict the statistical power of the distributions – the 95% confidence intervals from each distribution is shaded (integrated) in the other’s pdf. The statistical power BVM (denoted $P(A)$ in column B) is calculated for each model and indeed the model in row 2 is found to be preferable as it has the highest probability of agreement.

5.2 The $(\langle\epsilon\rangle, \beta_D)$ BVM

We invent a novel compound Boolean that defines agreement as when the model passes an average square error threshold of $\langle\epsilon\rangle$ and a check for probabilistic model representation. The later is imposed by requiring that $95\% \pm 4\%$ of the uncertain data lies inside the model’s $1 - \hat{\alpha} = 95\%$ confidence interval, i.e. $1 - \beta_D(\hat{\alpha}) \sim 95\%$. The $\pm 4\%$ tolerance was chosen such that overly uncertain models would be marked as “not agreeing” as they would be able to guarantee that 100% of the data lie within their overly wide confidence intervals. We call this compound Boolean the $(\langle\epsilon\rangle, \beta_D)$ Boolean. The BVM in this case is

$$p(A|M, D, \langle\epsilon\rangle, \beta_D) = \int_{\hat{Y}, Y} \rho(\hat{Y}|M, D) \cdot \Theta\left(B(\hat{Y}, Y, \langle\epsilon\rangle, \beta_D)\right) \cdot \rho(Y|D) d\hat{Y} dY, \quad (17)$$

where the compound Boolean $B(\hat{Y}, Y, \langle \epsilon \rangle, \beta_D)$ is equal to,

$$B\left(\frac{1}{N} \sum_i |y_i - \hat{y}_i| \leq \langle \epsilon \rangle\right) \wedge B\left(0.91 \leq \frac{1}{N} \sum_i \Theta(y_i \in [-c_{\hat{\alpha}}, c_{\hat{\alpha}}]_i) \leq 0.99\right), \quad (18)$$

N is the number of data points in $\{y_i\} = Y$, and $[-c_{\hat{\alpha}}, c_{\hat{\alpha}}]_i$ is the model's 95% confidence interval at comparison location x_i . We treat the model confidence intervals as certain quantities, which can be achieved effectively through enough Monte Carlo (MC) simulation of the model output pdfs (although this stipulation can be removed if needed).

Although the mathematical notation for the compound Boolean is a bit cumbersome, it is relatively easy to implement programmatically using *if* statements. This ease of programming allows the BVM to have a large capacity for representing complex and abstract validation scenarios in practice.

Expressing the BVM as an expectation value over $\rho(Y|D)\rho(\hat{Y}|M, D)$,

$$p(A|M, D, \langle \epsilon \rangle, \beta_D) = E\left[\Theta\left(B(\hat{Y}, Y, \langle \epsilon \rangle, \beta_D)\right)\right] \sim \frac{1}{K} \sum_{k=1}^K \Theta\left(B(\hat{Y}_k, Y_k, \langle \epsilon \rangle, \beta_D)\right), \quad (19)$$

allows one to compute the integral using standard statistical methods like MC. We use MC and $K = 3000$ samples in this toy example. In Figure 3 we implement the $(\langle \epsilon \rangle, \beta_D)$ Boolean and show that it is able to quantify both average error and a model's probabilistic representation of uncertain data, simultaneously.

We consider data that is generated from,

$$y(x) = 1 + x \exp(-\cos(10x)) + \sin(10x) + \epsilon_a(x), \quad (20)$$

where $\epsilon_a(x) \sim N(0, 0.4^2)$ represents aleatoric stochastic uncertainty due to the inherent randomness of the system. An instance of the aleatoric data Y without measurement uncertainty is depicted in red in Figure 3A) and B). We consider the data to have an additional epistemic measurement uncertainty $\epsilon_e(x) \sim N(0, 0.2^2)$ that contributes to the probability of whether or not the model agrees with the plotted instance of the aleatoric data Y .⁵

The models plotted in Figure 3A) and Figure 3B) are generated from,

$$\hat{y}(x; a, b, c, d, f, g) = a + bx \exp(-c \cos(dx)) + f \sin(gx), \quad (21)$$

where (a, b, c, d, f, g) are the model parameters. In Figure 3A), a deterministic model is considered and plotted in blue by treating the model parameters as completely certain numbers $(a, b, c, d, f, g) = (1, 1, 1, 10, 1, 10)$. In Figure 3B), we consider an uncertain model by treating the model parameters as uncertain values drawn from a multivariate Gaussian distribution having averages $(\mu_a, \mu_b, \mu_c, \mu_d, \mu_f, \mu_g) = (1, 1, 1, 10, 1, 10)$ and standard deviations $(\sigma_a, \sigma_b, \sigma_c, \sigma_d, \sigma_f, \sigma_g) = (0.35, 0.3, 0.3, 0.3, 0.3, 0.3)$. We evaluate the probability these data and model pairs pass the $\langle \epsilon \rangle$ Boolean verses the $(\langle \epsilon \rangle, \beta_D)$ Boolean by calculating their respective BVM's.

⁵Our framework is not limited to investigating other validations scenarios, e.g., if one has many aleatoric data path instances $\{Y\}$ (or enough statistics about the data pdf) that one would like to validate or use for modeling testing in the congregate.




Figure 3: The deterministic model plotted in Figure 3A) satisfies the average error validation requirement with $P(A|\langle\epsilon\rangle) = .99$ when a threshold of $\langle\epsilon\rangle = 0.46$. This result is logical because the congregate standard deviation of the data is itself the combination of the aleatoric and epistemic uncertainties is $\sqrt{0.4^2 + 0.2^2} \approx 0.45$. This deterministic model fails to predict the uncertain fluctuations of the data because deterministic models have confidence intervals with zero width. Thus the deterministic model fails to agree according to the $(\langle\epsilon\rangle, \beta_D)$ Boolean and one finds $P(A|\langle\epsilon\rangle, \beta_D) = 0$ for any value of $\langle\epsilon\rangle$. The uncertain model depicted in Figure 3B) is able to pass both agreement definitions; however, our choice to evaluate each probable model path against the epistemic uncertain data rather than just the average model increases the threshold to about $\langle\epsilon\rangle = 0.9$ before an agreement probability of about $P(A|\langle\epsilon\rangle) = .96$ is achieved. When instead evaluating $\langle\epsilon\rangle$ against the average model, we obtained similar results to the deterministic model for this Boolean $\langle\epsilon\rangle \sim 0.46$ and $P(A|\langle\epsilon\rangle) = .99$. The $(\langle\epsilon\rangle, \beta_D)$ BVM for the uncertain model is $P(A|\langle\epsilon\rangle, \beta_D) = 0.93$, because the uncertainty in the data more or less agrees with the confidence interval provided by the model. This model can be tested for agreement against other models or model parameter distributions for their respective definitions of agreement.

5.3 Exploring the BVM Ratio with the (γ, ϵ) BVM

In this subsection we invent an agreement function to represent the visual inspection an engineer might perform graphically and use the BVM ratio for model selection under this definition of agreement. By quantifying this, a practitioner could visually validate a model without actually looking at the model-data pair, which can be helpful for high dimensional spaces that are beyond human comprehension/visualizability. We will proceed by introducing this agreement function and some simple models to test it on. We will then quantify this measure using the BVM in the completely certain and uncertain cases. The main purpose of this example is to explore the BVM Ratio while showcasing the conceptual flexibility of the framework.

To quantify something resembling the visual inspection an engineer might make graphically, we use two main criteria. We define the model to be accepted if *most* of the model and data point pairs lie relatively close to one another *and* if none of the point pairs deviate too far from one another. We therefore consider a compound Boolean $B(\hat{Y}, Y, \gamma, \epsilon)$ that is true if a percentage larger than $\gamma\%$ ($\sim 90\%$) of the model output points \hat{Y} lie within ϵ of the data Y *and* 100% of the model output points lie within some multiple $m\epsilon$ of the data, which rules out obvious model form error. We will call this compound Boolean the (γ, ϵ) Boolean. The values $\gamma\%$, ϵ , and m can be adjusted to the needs of the modeler. It should be noted that the ϵ in this metric makes point by point evaluations as opposed to the average $\langle\epsilon\rangle$ Boolean used in the previous example. We will perform the analysis for a variety of γ and ϵ values to explore the limits of the metric.

We will calculate the BVM for two different order polynomial models that approximate N data points taken from the cosine function $y_i = \cos(x_i)$, as an illustration. The points are evenly spaced in the range $x_i \in [0, \pi]$. The first model $\hat{y}_i^{(1)} = M^{(1)}(x_i; a, b, c) = a + bx_i^2 + cx_i^4$ has uncertain parameters (a, b, c) and the second model $\hat{y}_i^{(2)} = M^{(2)}(x_i; a, b, c, d) = a + bx_i^2 + cx_i^4 + dx_i^6$ has uncertain parameters (a, b, c, d) .

To formulate the BVM, we still need to formulate the model and data probability distributions. Because

the Boolean expression $B(\hat{Y}, Y)$ is over the entire model and data functions, the model probability distribution is $p(\hat{Y}|M, D)$ and data probability distribution is $p(Y|D)$. These are joint probabilities over all of the points ($\hat{Y} = \{y_i\}$, $Y = \{y_i\}$) that constitute a particular path (\hat{Y}, Y) of the model or data, respectively. Because both models are linear in the uncertain coefficients (a, b, c, d), there is a one to one correspondence from the set of model parameters (a, b, c, d) to the set of the possible paths $\hat{Y}_{a,b,c,d}$ (given N is greater than the number of independent coefficients). This makes the uncertainty propagation from the uncertain model parameters to the full joint probability of the points on a path simple and results in the joint probability of paths being equal to the joint probability of the uncertain input model parameters. For simplicity we will let $p(a, b, c, d) = N(\mu_a, \sigma_a)N(\mu_b, \sigma_b)N(\mu_c, \sigma_c)N(\mu_d, \sigma_d)$, where $N(\mu, \sigma)$ is a normal distribution of average μ and standard deviation σ , such that is,

$$p(\hat{Y}_{a,b,c,d}|M_j) = \frac{1}{Z} \exp \left(-\frac{(a - \mu_a)^2}{2\sigma_a^2} - \frac{(b - \mu_b)^2}{2\sigma_b^2} - \frac{(c - \mu_c)^2}{2\sigma_c^2} - \frac{(d - \mu_d)^2}{2\sigma_d^2} \right).$$

Because the problem is well understood, we discretize the integrals rather than estimating them with MC. After discretization, the (γ, ϵ) BVM for each model is,

$$p(A|M_j, D, \gamma, \epsilon) = \sum_{\hat{Y}, Y} p(\hat{Y}|M_j) \cdot \Theta(B(\hat{Y}, Y, \gamma, \epsilon)) \cdot p(Y|D). \quad (22)$$

In principle $\epsilon = (\epsilon_1, \dots, \epsilon_N)$ is an N dimensional vector where each ϵ_i may be adjusted to impose more or less stringent agreement conditions on a point to point basis, which may be used to enforce reliability in regions of interest. In our example we let all of the components of ϵ_i be equal. If the standard deviation of each data point $\sim \hat{\sigma}_i$ (aleatoric and/or measurement uncertainty) in the joint data pdf $p(Y|D)$ is much less than ϵ_i and N is large, one may approximate the BVM as,

$$p(A|M_j, D, \gamma, \epsilon) \approx \sum_{\hat{Y}} p(\hat{Y}|M_j) \cdot \Theta(B(\hat{Y}, Y', \gamma, \epsilon)), \quad (23)$$

which can greatly reduce the number of combinations one must calculate by effectively treating the data as known, deterministic, and equal to Y' . We will use this approximation as it does not take away from the main point of this example.

We will use the following numerics. In the completely certain case, we will let the parameters be the Taylor series coefficients $(a, b, c, d) = (1, -\frac{1}{2!}, \frac{1}{4!}, -\frac{1}{6!})$ ($d = 0$ for model 1) and in the uncertain case we let each coefficient have Gaussian uncertainty centered at their Taylor series coefficients with standard deviations $(\sigma_a, \sigma_b, \sigma_c, \sigma_d) = (0.1, 0.05, 0.005, 0.0005)$ (and where $\sigma_d = 0$ for model 1). We let each model output path have $N = 50$ points and we allow for 20 possible values per parameter (a, b, c, d), which results in $20^3 = 8000$ possible paths for model 1 and $20^4 = 160,000$ for model 2. We let γ vary between 75% and 100% using an increment of 1% and let ϵ vary between 0 and 1 using an increment of 0.01. The value of m was chosen to be equal to 5, which imposes that no model path can have points that are greater than 5ϵ away while still be considered to agree with the data.

The BVM probability of agreement values as a function of the Boolean function parameters (γ, ϵ) are plotted in Figure 4 for model 1 and model 2 in the completely certain case:




Figure 4: Completely Certain Case: The BVM probability of agreement between model 1 and 2 with the data is plotted in the space of (γ, ϵ) . The results for model 1 (4th order polynomial) is plotted on the left and model 2 (6th order polynomial) is plotted on the right. Because here the models are deterministic, the BVM probability of agreements for each (γ, ϵ) pair is either zero or one. As expected, model 2 better fits the data in the space of (γ, ϵ) as it has more BVM values equal to one than model 1 as it is overall closer to the cosine function being that it is the next nonzero order in the Taylor series expansion. Neither model fits the data exactly as the BVM for both models at $(\gamma = 100\%, \epsilon = 0)$ is zero.

For a single (γ, ϵ) pair, the model's BVM ratio (a prior the models are assumed to be equally likely) is,

$$R(B(\gamma, \epsilon)) = \frac{p(A|M_1, D, \gamma, \epsilon)}{p(A|M_2, D, \gamma, \epsilon)}, \quad (24)$$

which, because the numerator and denominator is either 0 or 1 in the deterministic case, gives $R(B(\gamma, \epsilon))$ equal to 1, 0, ∞ , or $\frac{0}{0}$ meaning that the models both agree, model 1 does not agree but model 2 agrees, model 1 agrees but model 2 does not agree, or both models disagree, respectively. Thus, the BVM ratio for a single (γ, ϵ) pair between two deterministic models with completely certain data is not particularly insightful as they either agree or do not agree as defined by B . As it may not always be clear precisely what values of (γ, ϵ) one should choose to define agreement, one can meaningfully average (marginalize, analagous to (16)) over a viable volume in the space of (γ, ϵ) , with $p(\gamma, \epsilon) = \frac{1}{V}$, and arrive at an averaged Boolean BVM ratio,

$$R(B) = \frac{\sum_{\gamma, \epsilon} p(A|M_1, D, \gamma, \epsilon)}{\sum_{\gamma, \epsilon} p(A|M_2, D, \gamma, \epsilon)} = \frac{N_{1A}}{N_{2A}}, \quad (25)$$

which is simply a ratio of the number of agreements found for model 1, N_{1A} , in the (γ, ϵ) volume to the number of agreements found for model 2, N_{2A} , in the selected (γ, ϵ) volume. In our deterministic example, $R(B) = \frac{1108}{2364} = 0.4687$ as model 2 better fits the data, as defined by B , for the chosen meaningful (γ, ϵ) volume (which is taken to be the whole tested volume in this toy example). The BVM ratio or the averaged Boolean BVM ratio may be used as a guide for selecting models in the deterministic case given reasonable regions are chosen.

The BVM probability of agreement values as a function of (γ, ϵ) are plotted in Figure 5 for model 1 and model 2 in the uncertain model case:




Figure 5: Uncertain Case: The BVM probability of agreement between model 1 and 2 with the data is plotted in the space of (γ, ϵ) . The results for model 1 (4th order polynomial) is plotted on the left and model 2 (6th order polynomial) is plotted on the right. Because the model paths are uncertain, the BVM probability of agreements for each (γ, ϵ) pair may take any value from zero to one. As expected, model 2 better fits the data in the space of (γ, ϵ) as its BVM is generally larger than that of model 1; however, the BVM values are about equal in cases of large values of γ and ϵ values (the definition of agreement is less stringent and they both “agree”) and in the case of demanding absolute equality ($\epsilon = 0$) as neither model fits the data exactly.

The BVM ratios $R(B(\gamma, \epsilon)) = \frac{p(A|M_1, D, \gamma, \epsilon)}{p(A|M_2, D, \gamma, \epsilon)}$ of the uncertain models are plotted as a function of (γ, ϵ) in Figure 6:




Figure 6: This is a plot of the BVM ratios in the uncertain case. Model 2 is generally favored over model 1 as there exist no values greater than one on the plot. The amount the BVM ratio favors model 2 over model 1 decreases as the metric becomes less and less stringent (i.e. as γ decreases and ϵ increases). The $\epsilon = 0$ line was removed because neither model agrees with the data exactly.

The averaged Boolean BVM ratio for the uncertain models is,

$$R(B) = \frac{\sum_{\gamma, \epsilon} p(A|M_1, D, \gamma, \epsilon)}{\sum_{\gamma, \epsilon} p(A|M_2, D, \gamma, \epsilon)} = 0.7471, \quad (26)$$

which conforms to the notion that model 2 is, generally speaking, the preferable model, and which may be communicated with this single number. We given examples of the BVM ratio correctly selecting models according to abstract and new forms of validation.

6 Conclusion

We demonstrated the versatility of the BVM toward expressing and solving model validation problems. The BVM quantifies the probability the model is valid for arbitrary quantifiable definitions of model-data agreement using arbitrary quantified comparison functions of the model-data comparison values. The BVM was shown: to obey all of the desired validation metric criteria [34] (which is a first), to be able to represent all of the standard validation metrics as *special cases*, to supply improvements and generalizations to those special cases, and to be a tool for quantifying the validity of a model in novel model-data contexts. The later was demonstrated by the validation metrics we invented and quantified in our examples.

Finally, it was shown that one can perform model selection using the BVM ratio. The BVM model testing framework was shown to generalize the Bayesian model testing framework to arbitrary model-data contexts and with reference to arbitrary comparisons and agreement definitions. That is, the BVM ratio may be used to rank models directly in terms of the relevant model-data validation context. The problem of model-data validation may be reduced to the problem of finding/ defining the four BVM inputs: (\hat{z}, z) , $\rho(\hat{z}, z|M, D)$, $f(\hat{z}, z)$, and $B(f)$, and computing their BVM value. We find that the BVM is a useful tool for quantifying, expressing, and performing model validation and testing. Our future work involves probabilistically regressing (learning) model parameter distributions that satisfy arbitrary definitions of agreement within the BVM framework.

Acknowledgments This work was supported by the Center for Complex Engineering Systems (CCES) at King Abdulaziz City for Science and Technology (KACST) and the Massachusetts Institute of Technology (MIT). We would like to thank all of the researchers with the Center for Complex Engineering (CCES), especially Zeyad Al-Awwad, Arwa Alanqari, and Mohammad Alrished. Finally, we would also like to thank Nicholas Carrara and Ariel Caticha.

References

References

- [1] W. L. Oberkampf, T. G. Trucano, and C. Hirsch, 2004, Verification, Validation, and Predictive Capability in Computational Engineering and Physics, Appl. Mech. Rev., 57(3), pp. 345384.
- [2] D. Sornette, A. B. Davis, K. Ide, K. R. Vixie, V. Pisarenko, and J. R. Kamm, 2007, Algorithm for Model Validation: Theory and Applications, Proc. Natl. Acad. Sci. U.S.A., 104(16), pp. 65626567.
- [3] T.G. Trucano, L.P. Swiler, T. Igusa, W. L. Oberkampf, and M. Pilch, 2006, "Calibration, validation, and sensitivity analysis: Whats what", Reliab. Eng. Syst. Saf. 91, pp. 1331-1357.
- [4] M. C. Kennedy and A. O'Hagan, Bayesian calibration of computer models, J. R. Stat. Soc.: Ser. B (Stat Methodol.), vol. 63, no. 3, pp. 425-464, 2001.
- [5] O. P. Maître, O. M. Knio, Spectral Methods for Uncertainty Quantification. New York: Springer Science+Business Media, 2010.
- [6] Adams, B.M., Bauman, L.E., Bohnhoff, W.J., Dalbey, K.R., Ebeida, M.S., Eddy, J.P., Eldred, M.S., Hough, P.D., Hu, K.T., Jakeman, J.D., Stephens, J.A., Swiler, L.P., Vigil, D.M., and Wildey, T.M., "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.0 Users Manual," Sandia Technical Report SAND2014-4633, July 2014. Updated November 2015 (Version 6.3).

- [7] R. Ghanem, D. Higdon, and H. Owhadi, “The Uncertainty Quantification Toolkit (UQTk)” Handbook of Uncertainty Quantification, Springer, 2016.
- [8] B. Debusschere, N. Habib, N. Najm, P. Pébay, O. M. Knio, R. Ghanem, and O. P. Maître, Numerical Challenges in the Use of Polynomial Chaos Representations for Stochastic Processes, *SIAM J. Sci. Comput.* vol. 26, no. 2, pp. 698-719, 2005.
- [9] W. Gilks, S. Richardson, and D. J. Spiegelhalter, Markov chain Monte Carlo in practice. London, UK: Chapman and Hall, 1996.
- [10] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.*, vol. 21, no. 6, pp. 1087, 1953.
- [11] S. Sankararaman and S. Mahadevan, Integration of model verification, validation, and calibration for uncertainty quantification in engineering systems, *Reliab. Eng. Syst. Saf.*, vol. 138, pp. 194-209, 2015.
- [12] S. Sankararaman and S. Mahadevan, Model validation under epistemic uncertainty, *Reliab. Eng. Syst. Saf.* vol. 96, no. 9, pp. 1232-1241, 2011.
- [13] S. Mahadevan and R. Rebba, Validation of reliability computational models using Bayes networks, *Reliab. Eng. Syst. Saf.* vol. 87, no. 2, pp. 223-232, 2005.
- [14] C. Li and S. Mahadevan, Role of calibration, validation, and relevance in multi-level uncertainty integration, *Reliab. Eng. Syst. Saf.*, vol. 148, pp. 32-43, 2016.
- [15] S. Ferson, W. L. Oberkampf, and L. Ginzburg, Model validation and predictive capability for the thermal challenge problem, *Comput. Methods Appl. Mech. Eng.*, vol. 197, no. 29, pp. 2408-2430, 2008.
- [16] C. Roy and W. Oberkampf, A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing, *Comput. Methods Appl. Mech. Engng.*, vol. 200, pp. 2131-2144, 2011.
- [17] Y. Ling and S. Mahadevan, Quantitative model validation techniques: New insights, *Reliab. Eng. Syst. Saf.*, vol. 111, pp. 217, 2013.
- [18] W. Li, W. Chen, Z. Jiang, Z. Lu, and Yu Liu, New validation metrics for models with multiple correlated responses, *Reliab. Eng. Syst. Saf.*, vol. 127, pp. 1-11, 2014.
- [19] D. Wu, Z. Lun, Y. Wang, and L. Cheng, Model validation and calibration based on component functions of model output, *Reliab. Eng. Syst. Saf.*, vol. 140, pp. 59-70, 2015.
- [20] L. Zhao, Z. Lu, W. Yun, and W. Wang, Validation metric based on Mahalanobis distance for models with multiple correlated responses, *Reliab. Eng. Syst. Saf.*, vol. 159, pp. 80-89, 2017.
- [21] M. Stefano and B. Sudret, UQLab user manual - Polynomial chaos expansions. Report UQLab-V0.9-104, Chair of Risk, Safety and Uncertainty Quantification, ETH Zurich, 2015.
- [22] M. Parno and A. Davis, “MUQ: MIT Uncertainty Quantification Library”, <http://muq.mit.edu/home>, 2018.
- [23] C. Wang and H. G. Matthies, Novel model calibration method via non-probabilistic interval characterization and Bayesian theory, *Reliab. Eng. Syst. Saf.*, vol. 183, pp. 84-92, 2019.
- [24] R. Rebba and S. Mahadevan, Computational methods for model reliability assessment, *Reliab. Eng. Syst. Saf.*, vol. 93, no. 8, pp. 1197-1207, 2008.
- [25] N. T. Stevens, Assessment and comparison of Continuous Measurement Systems, Ph. D. Thesis, Dept. Statistics, Univ. of Waterloo, Waterloo, Ontario, Canada, 2014.

- [26] S. Sankararaman and S. Mahadevan, Assessing the reliability of computational models under uncertainty. In: The 54th AIAA/ASME/ASCE/AHS/ASC structures, structural dynamics, and materials conference; 2013.
- [27] W. L. Oberkampf, M. F. Barone, Measures of agreement between computation and experiment: validation metrics, *J. Comput. Phys.* vol. 217, no. 1, pp. 5-36, 2006.
- [28] W.L. Oberkampf, M.F. Barone, Measures of agreement between computation and experiment: validation metrics, *AIAA Paper* 2004 2626.
- [29] D. Sivia and J. Skilling, “Data Analysis A Bayesian Tutorial second edition”, Oxford University Press, Oxford, UK, 2006.
- [30] B. Placek, Bayesian Detection and Characterization of Extra-Solar Planets Via Photometric Variations, Ph. D. Thesis, Dept. Physics, Univ. at Alb. (S.U.N.Y.), Albany, NY, USA, 2014.
- [31] A. E. Gelfand and D. K. Dey, Bayesian model choice: asymptotics and exact calculations, *J. R. Stat. Soc. Ser. B (Methodol.)*, pp. 501-514, 1994
- [32] J. Geweke, Bayesian model comparison and validation, *Am. Econ. Rev.* 2007;97(2):604.
- [33] R. Zhang, S. Mahadevan S, Bayesian methodology for reliability model acceptance, *Reliab. Eng. Syst. Saf.*, vol. 80, no. 1, pp. 95-103, 2003.
- [34] Y. Liu, W. Chen, P. Arendt, and H. Z. Huang, Toward a better understanding of model validation metrics, *Trans ASME J. Mech. Des.* vol. 133, no. 7, pp. 071005, 2011.
- [35] K. A. Maupin, L. P. Swiler, and N. W. Porter, Validation metrics for deterministic and probabilistic data. *Journal of Verification, Validation and Uncertainty Quantification*, vol. 3, no. 3, pp. 031002, 2018.
- [36] G. Lee, W. Kim, H. Oh, B. D. Youn, and N. H. Kim, Review of statistical model calibration and validation from the perspective of uncertainty structures, *Structural and Multidisciplinary Optimization* 91, pp. 1-26, 2019.
- [37] J. Mullins, Y. Ling, S. Mahadevan, L. Sun, A. Strachan, Separation of aleatory and epistemic uncertainty in probabilistic model validation, *Reliab. Eng. Syst. Saf.*, vol. 147, pp. 49-59, 2016.
- [38] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine Learning in Python, *Journal of Machine Learning Research*, vol. 12, pp. 2825-2830, 2011.
- [39] E. T. Jaynes. In “Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science Vol. II”, D. Reidel Publishing Company, Dordrecht-Holland, pp. 175-257, 1976.
- [40] E. T. Jaynes, *Probability Theory: The Logic of Science*, UK Cambridge: Cambridge University Press, 2003.
- [41] F. Feroz and M. P. Hobson, Multimodal nested sampling: an efficient and robust alternative to Markov chain Monte Carlo methods for astronomical data analyses, *Monthly Notices of the Royal Astronomical Society*, vol. 384, no. 2, pp. 449-463, 2008.
- [42] A. Caticha, Entropic Inference and the Foundations of Physics (monograph commissioned by the 11th Brazilian Meeting on Bayesian Statistics - EBEB-2012). 2012. URL <http://www.albany.edu/physics/ACaticha-EIFP-book.pdf>.
- [43] K. Knuth, Optimal Data-Based Binning for Histograms, [arxiv:physics/0605197v2](https://arxiv.org/abs/physics/0605197v2), <https://arxiv.org/abs/physics/0605197v2>, 2013.

A Deriving the other validation metrics from the BVM

In the following subsections we will show some of the special cases of the Bayesian validation metric. Subsequent improvements or immediate generalizations of the metrics using (2) are presented when applicable. A detailed review of the majority of these metrics may be found in [34] and the references therein. Table 1 and 2 in Section 4 outline the results.

A.1 Reliability metric and probability of agreement

There are a few validation metrics related to the reliability metric present in the literature. The reliability metric $r = p(|\langle \hat{y} \rangle - \mu_y| < \epsilon)$ [24] is equal to the probability that the data and the model expectation values are within a tolerance of size ϵ . Their “probability of agreement” introduced in [25] is closely related to r , but instead expresses the quantity as “the probability the data and the model expectation values *agree* within a tolerance (or sliding tolerance) of ϵ ”. The reliability metric was expanded in [26] to account for model outputs and data rather than simply comparing the mean of the model prediction against the mean of the data. The improved reliability metric is equal to,

$$r_i = \int_{-\infty}^{\infty} \rho(Y|D) \int_{Y-\epsilon(Y)}^{Y+\epsilon(Y)} \rho(\hat{Y}|M) d\hat{Y} dY, \quad (27)$$

where $\rho(\hat{Y}|M)$ and $\rho(Y|D)$ are the full joint probability distributions of the model outputs and the data, respectively. This metric quantifies the probability that the error is less than a value $\epsilon(y)$ on a point to point basis.

The BVM is the reliability metric when the comparison values are $\hat{z} = \langle \hat{y} \rangle$ and $z = \mu_y$, and B takes the form of an inequality, being true if $-\epsilon \leq \langle \hat{y} \rangle - \mu_y \leq \epsilon$. If we would like to use a sliding interval of “tolerance” or “error acceptance”, denote it by $[c_-(\mu_y), c_+(\mu_y)]$, where $c_- < c_+$, and the BVM is,

$$\begin{aligned} p(A|M, D) &= \int_{\langle \hat{y} \rangle, \mu_y} \rho(\langle \hat{y} \rangle|M) \Theta\left(c_-(\mu_y) \leq \langle \hat{y} \rangle \leq c_+(\mu_y)\right) \rho(\mu_y|D) d\langle \hat{y} \rangle d\mu_y \\ &= \int_{\mu_y} \int_{\langle \hat{y} \rangle=c_-(\mu_y)}^{c_+(\mu_y)} \rho(\langle \hat{y} \rangle|M) \rho(\mu_y|D) d\langle \hat{y} \rangle d\mu_y = r. \end{aligned} \quad (28)$$

This is the reliability metric if $c_{\pm} = \mu_y \pm \epsilon$ is a constant and symmetric interval about μ_y .

The BVM is the improved reliability metric when $\hat{z} = \hat{Y}$, $z = Y$, and when the Boolean $B(\hat{z}, z)$ is true iff $|\hat{y} - y| \leq \epsilon(y)$ for all \hat{y}, y pairs. This is,

$$p(A|M, D) = \int_{\hat{Y}, Y} \rho(Y|D) \rho(\hat{Y}|M) \left(\prod_{\hat{y}, y \text{ pairs}} \Theta(|\hat{y} - y| \leq \epsilon(y)) \right) d\hat{Y} dY = r_i. \quad (29)$$

The BVM quantifies the probability of square error (or difference) is less than some ϵ by considering,

$$p(A|M, D) = \int_{\hat{Y}, Y} \rho(Y|D) \Theta(|\hat{Y} - Y| \leq \epsilon) \rho(\hat{Y}|M) d\hat{Y} dY. \quad (30)$$

Nothing in the BVM requires the variables to be continuous or ordered, so the natural generalization is to let the Boolean expression be true if “The value \hat{z} is in the subset $S(z)$, which is the set of \hat{z} ’s agreeing with z ”. For example, if \hat{z} ’s are strings, $S(z)$ might be the set of words or phrases in \hat{z} that are reasonably synonymous with z . This gives the straightforward generalization to accommodate arbitrary data types,

$$p(A|M, D) = \sum_{\hat{z}, z} p(\hat{z}|M) \Theta\left(\hat{z} \in S(z)\right) p(z|D) = \sum_z \sum_{\hat{z} \in S(z)} p(\hat{z}|M) p(z|D), \quad (31)$$

by using sets rather than intervals.

A.2 Frequentist validation metric

To include the frequentist validation metric in the BVM, we will have to express the comparison variables \hat{z} and z and their respective probabilities. The result can be replicated by letting: $z = \mu_y$ be the Student-t distribution and $\hat{z} = \langle \hat{y} \rangle$ have a Dirac delta distribution $\rho(\hat{z}|M, D) = \rho(\langle \hat{y} \rangle|M, D) = \delta(\langle \hat{y} \rangle - \langle \hat{y} \rangle')$ where $\langle \hat{y} \rangle'$ is the known value of the computational model's expected output. Because the frequentist validation metric does not force the modeler to define what is meant by *agreement*, we represent this freedom by keeping $B(\hat{z}, z)$ general. This gives,

$$\begin{aligned} p(A|M, D) &= \int_{\hat{z}, z} \rho(\hat{z}|M, D) \Theta(B) \rho(z|D) d\hat{z} dz \\ &= \int_{\langle \hat{y} \rangle, \mu_y} \delta(\langle \hat{y} \rangle - \langle \hat{y} \rangle') \cdot \Theta(B(\langle \hat{y} \rangle, \mu_y)) \\ &\quad \cdot \left(\frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi}\Gamma(\frac{\nu}{2})} \left(1 + \frac{(\mu_y - \bar{y})^2}{\nu \bar{s}^2 / N} \right)^{-\frac{\nu+1}{2}} \right) d\langle \hat{y} \rangle d\mu_y, \end{aligned} \quad (32)$$

where \bar{y} is the population average, \bar{s} is the population standard deviation, and ν is the degrees of freedom are the parameters of the data's Student-t distribution. Making the coordinate transformations $\langle \hat{y} \rangle \rightarrow \bar{E} = \langle \hat{y} \rangle - \bar{y}$ and $\mu_y \rightarrow E = \langle \hat{y} \rangle - \mu_y$ gives,

$$\begin{aligned} p(A|M, D) &= \int_{\bar{E}, E} \delta(\bar{E} - \bar{E}') \cdot \Theta(B(\bar{E}, E)) \cdot \left(\frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi}\Gamma(\frac{\nu}{2})} \left(1 + \frac{(E - \bar{E})^2}{\nu \bar{s}^2 / N} \right)^{-\frac{\nu+1}{2}} \right) d\bar{E} dE \\ &= \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi}\Gamma(\frac{\nu}{2})} \int_{E'} \Theta(B(\bar{E}', E')) \cdot \left(1 + \frac{(E' - \bar{E}')^2}{\nu \bar{s}^2 / N} \right)^{-\frac{\nu+1}{2}} dE', \end{aligned} \quad (33)$$

with $E' = \langle \hat{y} \rangle' - \mu_y$ and $\bar{E}' = \langle \hat{y} \rangle' - \bar{y}$. Given that judgments of agreement in the frequentist validation metric are expected to be made based on the confidence level $1 - \alpha$ that E' is within the confidence interval, this may be factored into $B(\bar{E}', E') \rightarrow B(\bar{E}', E', \alpha)$, as well as other user defined terms toward expressing agreement. Equation (33) is thought to be the full BVM's representation of the frequentist validation metric.

The immediate generalization to the frequentist validation metric offered by the BVM is to let the model output expectation value $\hat{z} = \langle \hat{y} \rangle$ have some amount of uncertainty. The uncertainty is perhaps Gaussian or Student t distributed in $\mu_{\langle \hat{y} \rangle}$ (the true model expectation value) due to only having a finite number of Monte Carlo samples and/or uncertainty induced by discretization error. Because $\hat{z} = \mu_{\langle \hat{y} \rangle}$ and $z = \mu_y$ are both uncertain in general, one generalizes to,

$$p(A|M, D, B) = \int_{\mu_{\langle \hat{y} \rangle}, \mu_y} \rho(\mu_{\langle \hat{y} \rangle}|M) \cdot \Theta(B(\mu_{\langle \hat{y} \rangle}, \mu_y)) \cdot \rho(\mu_y|D) d\mu_{\langle \hat{y} \rangle} d\mu_y. \quad (34)$$

It is interesting to note the consequence of defining a reasonable Boolean expression of agreement on the BVM representation of the frequentist metric. A natural agreement function B for the metric is one that is true if $E_\mu = |\mu_{\langle \hat{y} \rangle} - \mu_y| \leq \epsilon$. The BVM then gives,

$$\begin{aligned} p(A|M, D) &= \int_{\mu_{\langle \hat{y} \rangle}, \mu_y} p(\mu_{\langle \hat{y} \rangle}|M, D) \cdot \Theta(E_\mu \leq \epsilon) \cdot p(\mu_y|D) d\mu_{\langle \hat{y} \rangle} d\mu_y \\ &= \int_{\mu_{\langle \hat{y} \rangle}} \int_{\mu_y = \mu_{\langle \hat{y} \rangle} \pm \epsilon} p(\mu_{\langle \hat{y} \rangle}|M, D) p(\mu_y|D) d\mu_{\langle \hat{y} \rangle} d\mu_y = r. \end{aligned} \quad (35)$$

Thus, the frequentist's validation metric is the reliability metric [24] and "probability of agreement" [25] when reasonable accuracy requirements are imposed on the acceptable difference between the expectation values.

A.3 Area and Binned Probability Difference Metric

The BVM is able to represent and generalize the area metric by letting the comparison values (z, \hat{z}) be the the cdfs in question to be compared $(F(y|D), F(\hat{y}|M))$. The Area metric is,

$$d[F(\hat{y}|M), F(y|D)] = \int_{-\infty}^{\infty} \left| F(\hat{y}|M) - F(y = \hat{y}|D) \right| d\hat{y}. \quad (36)$$

The Area metric may be represented by the BVM in a simple way. First allow the comparison value function to be the Area metric functional,

$$f(z, \hat{z}) \equiv d[F(\hat{y}|M), F(y|D)], \quad (37)$$

and define agreement with a Boolean,

$$p(A|M, D) = \Theta \left(B \left(d[F(\hat{y}|M), F(y|D)] \right) \right). \quad (38)$$

Here, the cdfs are treated as completely certain ($\delta(\hat{z} - F(\hat{y}|M))$ and $\delta(\hat{z} - F(\hat{y}|M))$ are Dirac delta functionals). The functional form of the Boolean may be decided on given the user's specific validation requirements; however, satisfying some kind of ϵ threshold $d[F(\hat{y}|M), F(y|D)] \leq \epsilon$ seems logical. Generalizations to the Area metric may be represented analogously with the BVM.

When the values of the cdfs are uncertain the BVM becomes,

$$p(A|M, D) \leftarrow \int \Theta \left(B \left(d[F(\hat{y}|M), F(y|D)] \right) \right) \rho(F(\hat{y}|M), F(y|D)) dF(\hat{y}|M) dF(y|D), \quad (39)$$

which may pose computational challenges and require random sampling; however, in some cases it may be permissable to treat the model cdf as known. The extension to uncertain cdf's is usually not considered in the literature; however, the theoretical generalization to the uncertain case is apparent in the BVM framework. As an approximation, one may consider discretizing the Area metric by breaking it into K comparison points,

$$d[F(\hat{y}|M), F(y = \hat{y}|D)] \approx \sum_{i=1}^K \left| F(y_i|M) - F(y = y_i|D) \right|. \quad (40)$$

The uncertainty in the cumulative distribution of the data $\rho(z|D)$ is now a finite joint product pdf $\rho(F(y_1|D), \dots, F(y_K|D)|D)$ over possible cdf values in the K bins, constrained by $F(y_i|D) \leq F(y_{i+1}|D)$. This metric may also be represented with the BVM.

Alternatively, one may consider a binned probability difference comparison functional,

$$d_m[p(\hat{y}|M), p(y = \hat{y}|D)] = \sum_{i=1}^K \left| p(y_i|M) - p(y = y_i|D) \right|. \quad (41)$$

Let

$$p(z|D) = P(p_{y_1}, \dots, p_{y_K}|D, \sum_k p_{y_k} = 1)$$

be the uncertainty for the probability estimate in each data bin due to the random process. Given that the pdf of the model is set to a (presumably known) single pdf function $p(\hat{y}|M)$, the BVM is,

$$\begin{aligned} p(A|M, D) &= \int_{\hat{z}, z} \delta(\hat{z} - \{p(\hat{y}|M)\}) \Theta \left(B \left(d_m[\hat{z}, z] \right) \right) p(z|D) d\hat{z} dz \\ &= \int_{p_{y_1}, \dots, p_{y_K}} \Theta \left(B \left(d_m[\{p(\hat{y}|M)\}, \{p(y_i|D)\}] \right) \right) P(p_{y_1}, \dots, p_{y_K}|D) \prod_i (dp_{y_i}). \end{aligned} \quad (42)$$

Given the form of the distribution $P(p_{y_1}, \dots, p_{y_K} | D)$ is well known (i.e. a Dirichlet distribution after N independent observations of y), it can be treated as an expectation value,

$$p(A|M, D) = E \left[\Theta \left(B \left(d_m[p(\hat{y}|M), \{p(y_i|D)\}] \right) \right) \right]_{\{p(y_i|D)\}}, \quad (43)$$

and estimated with Monte Carlo. When K is large we may face the curse of dimensionality if the probabilities (bin heights) are themselves uncertain.

The number of bins K plays a role similar to ϵ in that its choice affects the definition/context of agreement represented by the BVM. The binned pdf metric is most informative when K is large because one is checking each region of the pdf for agreement. On the other hand, perhaps when data is limited, there may be instances when having $K = 2$ bins is useful, if for instance one was interested in testing a model for representing binary type probabilities (like pass/fail or positive/negative). Using a favorable agreement in this simple binary context to imply that the model works for *otherwise untested* $K > 2$ comparisons/contexts/validations, is statistical misrepresentation and should be avoided (see Section 2.4). The BVM requires one to explicitly state the comparison values, the comparison value function, and the definition of agreement to compute $P(A|M, D)$ such that confusion may be avoided and model validation comparisons may be justified.

Thus, the choice of K ultimately determines the granularity of the definition of agreement being tested. If one uses methods similar to [43] to select K , then one is letting the complexity of the data and the number of data points determine the stringency of the definition of agreement.

A.4 Probability Density Function Comparison Metrics

Another way to gauge the agreement between uncertain data and models is through a pdf comparison metric, $G(\rho_D || \rho_M)$ [35]. Examples of $G(\rho_D || \rho_M)$ are the negative relative entropy or KL divergence,

$$D(\rho_D || \rho_M) = \int_y \rho(y|D) \log \left(\frac{\rho(y|D)}{\rho(\hat{y} = y|M)} \right) dy, \quad (44)$$

the Symmetrized KL divergence $S(\rho_D || \rho_M) = D(\rho_D || \rho_M) + D(\rho_M || \rho_D)$, the Jensen-Shannon divergence $D_{JS}(\rho_D || \rho_M)$, the Hellinger Metric $H(\rho_D || \rho_M)$, the Fisher information distance $\ell(\rho_D || \rho_M)$, and the Wasserstein distance $W(\rho_D || \rho_M)$. These metrics give a notion of “closeness” between the pdfs that can be used for validation.

The BVM may represent these validation metrics by letting the comparison value function $f(z, \hat{z})$ be the pdf comparison metric,

$$f(z, \hat{z}) \equiv G(\rho_D || \rho_M),$$

where $z \equiv \rho_M$ and $\hat{z} \equiv \rho_D$ and are the model and data pdfs, respectively. In the absence of model and data uncertainty (i.e. the pdfs are treated as known functions, e.g. ρ_D is a gaussian with mean=0 and variance=1), the BVM is simply,

$$p(A|M, D) = \Theta \left(B[f(z, \hat{z})] \right) = \Theta \left(B[G(\rho_D || \rho_M)] \right),$$

which either meets the specifications for them to agree defined by B , or does not (e.g. passing a tolerance threshold). Following the structure of the BVM, if there are uncertainties in the functional forms of the pdf’s, they may be included into the BVM,

$$p(A|M, D) \leftarrow \int_{\rho_D, \rho_M} \Theta \left(B[G(\rho_D || \rho_M)] \right) \rho(\rho_D, \rho_M) d\rho_D d\rho_M.$$

Uncertainties in the data pdfs may come from a lack of data and uncertainties in the model may come from parametric uncertainty (e.g. a Gaussian pdf model $\rho_{M|\mu}$ with an uncertain mean $\rho(\mu)$)⁶ Similar to the Area metric, these metrics may be discretized from pdf to probability comparison metrics and uncertainty in the pdfs themselves are typically not discussed/quantified in the literature.

⁶ One should note that here there is the potential for two different models. One in which $\rho_M \equiv \int \rho_{M|\mu} \rho(\mu) d\mu$ is marginalized over (in which case the model pdf is “certain” if integrated analytically) and another in which the model pdf is gaussian $\rho_{M|\mu}$ but there is model parametric uncertainty of the form $\rho(\mu)$.

A.5 Statistical hypothesis testing

Normally when statistical hypothesis testing is performed, one constructs the pdf of the relevant test statistic of the data $\rho(z|D)$ and then assumes the null hypothesis is true ($M = D$) counterfactually, which is enforced by setting the “to be tested population” (the model outputs here) pdf to be equal to the pdf of the test statistic of the data $\rho(\hat{z}|M) \rightarrow \rho(\hat{z}|M = D)$. However, in the present case, we are interested in the general modeling case in which one is able to extract a pdf of the outputs of the model, which we would like to test against data before assuming that they are equal. We will first represent the classical statistical hypothesis test using the BVM and then later supply a version that is more relevant to model validation problems.

Classical statistical hypothesis testing In classical hypothesis testing, one constructs the pdf of a relevant test statistic $S_y \equiv z$ of the data $\rho(z|D)$. The null hypothesis is that the model is equal to the data, which is enforced by setting $\rho(\hat{z}|M) = \rho(\hat{z}|M = D)$. Further, the null hypothesis is not rejected if the test statistic from the model \hat{z} falls within the critical region $[-c_\alpha, c_\alpha]$ that corresponds to the probability $\int_{-c_\alpha}^{c_\alpha} \rho(z|D) dz = 1 - \alpha$ of the data. The case of not rejecting the null hypothesis is represented by the Boolean expression $B(\hat{z})$, which is true if $-c_\alpha \leq \hat{z} \leq c_\alpha$ is true – defining “agreement” in this case. This results in the following BVM,

$$\begin{aligned} p(A|M = D, D) &= \int_{\hat{z}, z} \rho(\hat{z}|M = D) \Theta(B(\hat{z})) \rho(z|D) d\hat{z} dz \\ &= \int_{\hat{z}} \rho(\hat{z}|M = D) \Theta(-c_\alpha \leq \hat{z} \leq c_\alpha) d\hat{z} = 1 - \alpha, \end{aligned} \quad (45)$$

because the model was assumed to be equal to the data counterfactually in the null hypothesis and $B(\hat{z}, z) = B(\hat{z})$ only.

The probability of type I error, i.e. rejecting the counterfactually assumed true null hypothesis when it is actually true, is equal to α . It should be noted that this is not equal to the probability of finding the model value outside of the data’s confidence interval because it is unknown if the null hypothesis (in what results in $\rho(\hat{z}|M) \rightarrow \rho(\hat{z}|M = D)$) is actually true, or not, because the null hypothesis was merely assumed to be true counterfactually. It should further be noted that with probability α the data’s test statistic is outside of its *own* confidence interval. Thus, the probability α both indicates type I error *and* a systematic type error that the wrong sort of comparison is being made, i.e. the wrong Boolean expression was chosen, because, why would we care if the model is within a certain confidence interval if the data is not even within that interval?⁷ The probability of type II error, i.e. that the null hypothesis was not rejected when it is actually false, is equal to $\beta_M(\alpha) = 1 - \int_{-c_\alpha}^{c_\alpha} \rho(\hat{z}|M) d\hat{z}$, which in the classical case is difficult to calculate directly because one does not have access to the actual model pdf $p(\hat{z}|M)$ in frequentist probability.

Improved statistical hypothesis testing for validation For the validation cases we are interested in, both $\rho(\hat{z}|M)$ and $\rho(z|D)$ are quantified, and therefore assuming that $\rho(\hat{z}|M) \rightarrow \rho(\hat{z}|M = D)$ would irresponsibly throw away any information sent through the model. We therefore offer the improved statistical hypothesis test for validation using the BVM, which uses both the model and data pdfs. We call this BVM the statistical power BVM.

For the modified statistical hypothesis test, let the definition of *agreement* be a compound Boolean expression that is true iff both $-c_\alpha \leq \hat{z} \leq c_\alpha$, that the model test statistic lies in the data’s confidence interval, *and* $-c_{\hat{\alpha}} \leq z \leq c_{\hat{\alpha}}$ that the data statistic lies in the model’s confidence interval, which corresponds to the probability $\int_{-c_{\hat{\alpha}}}^{c_{\hat{\alpha}}} \rho(\hat{z}|M) d\hat{z} = 1 - \hat{\alpha}$ of the model. By not assuming $M = D$, we remove the possibility that either type I or type II errors can occur; however, systematic errors, that the Boolean expression meant to define agreement is nonsensical, still exist. Thus, while more or less adhering to the type of tests one

⁷Independent of whether or not the null hypothesis is true*

might perform for model validation using statistical hypothesis testing, we get,

$$\begin{aligned} p(A|M, D) &= \int_{\hat{z}, z} \rho(\hat{z}|M, D) \Theta(-c_\alpha \leq \hat{z} \leq c_\alpha) \Theta(-c_{\hat{\alpha}} \leq z \leq c_{\hat{\alpha}}) \rho(z|D) d\hat{z} dz \\ &= (1 - \beta_M(\alpha)) \cdot (1 - \beta_D(\hat{\alpha})), \end{aligned} \quad (46)$$

which is the probability that both the model and the data lie within one another's confidence intervals. The value $1 - \beta(\alpha)$ is called the statistical power of the test, but here we have access to both the statistical power of the data and the model. The probability the model and data do not agree as defined by B is given by,

$$p(\bar{A}|M, D) = 1 - p(A|M, D) = \beta_D(\hat{\alpha}) + \beta_M(\alpha) - \beta_D(\hat{\alpha})\beta_M(\alpha), \quad (47)$$

which occurs if either \hat{z} , z , or both are outside of one another's confidence intervals. The probability for systematic error (that \hat{z} , z , or both are outside of their own confidence intervals) is equal to $\alpha + \hat{\alpha} - \alpha\hat{\alpha}$; however, there is no conceptual issue with setting α , $\hat{\alpha}$, or both equal to 0 as long as both distributions do not span the entire range of possible values (as in such a case the model and data would always agree, which means the test has zero resolving power). Setting $\alpha = \hat{\alpha} = 0$ removes the chance of systematic error from the analysis.

Overall, the use of confidence intervals is suboptimal unless both the model and data pdfs are strictly unimodal. Therefore, rather than using a confidence interval, one may use a "confidence set", which we define as the smallest set of \hat{z} (as well as for z) values that's probability adds to $1 - \hat{\alpha}$ (and similarly $1 - \alpha$). Confidence sets are generated by adding the largest probabilities of \hat{z} (z) until a confidence level of $1 - \hat{\alpha}$ ($1 - \alpha$) is met. A confidence interval is only equal to the confidence set if the distribution is unimodal. As the confidence set is the smallest set of values adding up the the confidence level, this set of values is more informative than a confidence interval, which may include many 0 or low probability events. Using a confidence set improves the resolving power of the metric further.

The statistical power BVM (46) is more informative than the classical statistical hypothesis test (45) because it utilizes the model pdf while also removing type I, II, and (optionally) systematic errors from the test; however its overall resolving power is weak when compared to other metrics. By rewriting the pair of overlapping confidence intervals (or sets) as the set of values in a single "overlap interval" $I = [c_\alpha, c_\alpha] \cap [c_{\hat{\alpha}}, c_{\hat{\alpha}}]$, one may see that (46) is a particular case of (31) with $S(y) = I$ for $y \in I$ and being the null set otherwise. Thus, the statistical power BVM may be seen as a special case of the generalized reliability metric suggested by the BVM in (31) that effectively has large and unvarying tolerance intervals (sets). The use of confidence intervals as well as confidence sets as a means to define agreement effectively coarse grain the probabilities, which removes their informative features. The most stringent definition of agreement between the model and the data leads to Bayesian model testing, which we prove in the next section, as it is indeed equal to the generalized model reliability metric with a zero tolerance for differences $\epsilon = 0$.

A.6 Bayesian model testing

In Bayesian model testing, rather than assuming a particular model is true, one lets the available data determine which model is most likely given that data. Represented probabilistically, out of a set of possible models, Bayesian model testing selects the model with the maximum posterior probability $p(M|Y)$, i.e. the probability of a (previously calibrated [11]) model M given the data set $Y = \{y_i\}$ from data source D . As the selection rule for models requires comparing values of $p(M|Y)$, one often constructs the posterior odds ratio,

$$R = \frac{p(M|Y)}{p(M'|Y)}, \quad (48)$$

and selects the model M with the highest value of R relative to a chosen base model M' . Using Bayes Theorem, the posterior odds ratio may be recast as,

$$R = \frac{\rho(Y|M)p(M)}{\rho(Y|M')p(M')} = K \frac{p(M)}{p(M')}, \quad (49)$$

where $K \equiv \frac{\rho(Y|M)}{\rho(Y|M')}$ is known as the Bayes factor, which is the ratio of the likelihoods of the models. The ratio of the prior model probabilities $\frac{p(M)}{p(M')}$ is the ratio of the belief that model M is true prior to examining the data, relative to M' . If there is reason to believe that one model is more probable than another due to prior (perhaps statistical) knowledge of the system under investigation, then the Bayesian framework allows one to take this into account through the prior model probabilities. As it is common for the R values to be potentially many orders of magnitude greater than one, the prior model probabilities may be overwhelmed by the Bayes factor. In any case, if there is no reason to prefer one mode over another, one can let the prior model probabilities be equal a-priori.

Given we are testing some model function $\hat{y} = M(\vec{x}, \vec{\alpha})$, we have access to the forward propagation of data and parameters through a given model, the problem of calculating R may be rerouted to the computation of K . Using the potentially multidimensional inputs to our models $(\vec{x}, \vec{\alpha})$, which represent the input data and the model parameters respectively, the Bayes factor is calculated through forward propagation (marginalization) of these inputs through both models,

$$K = \frac{p(Y|M)}{p(Y|M')} = \frac{\int_{\vec{x}, \vec{\alpha}} \rho(Y|\vec{x}, \vec{\alpha}, M) \rho(\vec{x}, \vec{\alpha}|M) d\vec{x} d\vec{\alpha}}{\int_{\vec{x}', \vec{\alpha}'} \rho(Y|\vec{x}', \vec{\alpha}', M') \rho(\vec{x}', \vec{\alpha}'|M') d\vec{x}' d\vec{\alpha}'} \quad (50)$$

The prior probability to the inputs of the model $\rho(x, \alpha|M)$ are the input probability distributions used to propagate uncertainty through the model. The probability $\rho(Y|x, \alpha, M)$ is the probability of the data given by the knowledge of the model function $\hat{y} = M(x, \alpha)$; however, it should be noted that in general the data $Y \neq \hat{Y} = \{y_i\}$, is not the set of model outputs \hat{Y} , as the data Y was collected from the experiment rather than from model outputs. Thus, one must impose the assumptions under which a model is built, i.e. it is built for the purpose of approximating $y_i \approx \hat{y}_i = M(x_i, \alpha)$, as we will see later. Thus, a more verbose presentation of K is,

$$K = \frac{p(\hat{Y} \equiv Y|M)}{p(\hat{Y} \equiv Y|M')} = \frac{\rho(\hat{Y} \equiv Y|M)}{\rho(\hat{Y} \equiv Y|M')} = \frac{\int_{\vec{x}, \vec{\alpha}} \rho(\hat{Y} = Y|\vec{x}, \vec{\alpha}, M) \rho(\vec{x}, \vec{\alpha}|M) d\vec{x} d\vec{\alpha}}{\int_{\vec{x}', \vec{\alpha}'} \rho(\hat{Y} = Y|\vec{x}', \vec{\alpha}', M') \rho(\vec{x}', \vec{\alpha}'|M') d\vec{x}' d\vec{\alpha}'} \quad (51)$$

where $p(\hat{Y} \equiv Y|M)$ is understood to be the sum of the model and the data probabilities that jointly output the same values, i.e. the probability that the any of the possible model and data values turn out to be the same. The form of $\rho(\hat{Y} = Y|\vec{x}, \vec{\alpha}, M)$ is usually assumed to be something that works (and even better if it is also simple) [30], like the product of Gaussian distributions,

$$\rho(\hat{Y} = Y|\vec{x}, \vec{\alpha}, M) = \frac{1}{Z} \exp\left(-\frac{1}{2\sigma^2} \sum_i (M(x_i; \vec{\alpha}) - y_i)^2\right), \quad (52)$$

where σ^2 is interpreted as the measurement uncertainty of the data.⁸ One usually computes the integrals in K using various sampling algorithms such as nested sampling [29, 41] or another Markov Chain Monte Carlo technique. As an added bonus, Bayesian model testing has an inbuilt Occam's razor mechanism which penalizes needlessly complex models, i.e. one that would over fit the data by using a large number of uncertain model parameters $\vec{\alpha}$. A clear explanation may be found in [42].

We can represent Bayesian model testing using the Bayesian validation metric. Let the model comparison value \hat{z} be $\hat{Y} = (y_1, \dots, y_N)$ that in principle correspond to $z = Y = (y_1, \dots, y_N)$, a validation set of data. The Boolean expression B is considered to factor into a set of "and" statements over the individual model output and data points $B(\hat{Y}, Y) = B(y_1, y_1) \wedge \dots \wedge B(y_N, y_N)$, where each B is true iff $y_i = \hat{y}_i$ exactly. The Bayesian validation metric in this case is then,

$$\begin{aligned} p(A|M, D) &= \int_{\hat{Y}, Y} \rho(\hat{Y}|M, D) \theta(B(\hat{Y}, Y)) \rho(Y|D) d\hat{Y} dY \\ &= \int_{\hat{Y}, Y} \rho(\hat{Y}|M, D) \delta_{\hat{Y}, Y} \rho(Y|D) d\hat{Y} dY. \end{aligned} \quad (53)$$

⁸The dependence on the data set/ experiment D is therefore implied.

In general, computational models may be described by a model function $\hat{Y} = \vec{y} = M(\vec{x}, \vec{\alpha})$, and given the model pdf was constructed through forward propagation of the uncertainties $\rho(\vec{x}, \vec{\alpha})$, the model output pdf is

$$\rho(\hat{Y}|M, D) = \int_{\vec{x}, \vec{\alpha}} \rho(\hat{Y}|\vec{x}, \vec{\alpha}, M, D) \rho(\vec{x}, \vec{\alpha}) d\vec{x} d\vec{\alpha} \quad (54)$$

Substituting (54) into (53), using the trick (11) and (12) – that

$$\int_{Y-\epsilon}^{Y+\epsilon} \rho(\hat{Y}|M, D) d\hat{Y} \xrightarrow{\epsilon \rightarrow 0^+} p(\hat{Y} = Y|M, D) = \rho(\hat{Y} = Y|M, D) d\hat{Y},$$

and integrating over Y , one finds,

$$\begin{aligned} p(A|M, D) &= \int_{\vec{x}, \vec{\alpha}, Y} p(\hat{Y} = Y|\vec{x}, \vec{\alpha}, M, D) \rho(\vec{x}, \vec{\alpha}) \rho(Y|D) d\vec{x} d\vec{\alpha} dY \\ &= \int_{\vec{x}, \vec{\alpha}} p(\hat{Y} \equiv Y|\vec{x}, \vec{\alpha}, M, D) \rho(\vec{x}, \vec{\alpha}) d\vec{x} d\vec{\alpha} \\ &= \left(\int_{\vec{x}, \vec{\alpha}} \rho(\hat{Y} \equiv Y|\vec{x}, \vec{\alpha}, M, D) \rho(\vec{x}, \vec{\alpha}) d\vec{x} d\vec{\alpha} \right) d\hat{Y} \\ &= \rho(\hat{Y} \equiv Y|M, D) d\hat{Y}, \end{aligned} \quad (55)$$

which is what is meant by, and is equal to, the probability in the numerator of the Bayes factor (51),

$$p(A|M, D) = \rho(\hat{Y} \equiv Y|M, D) d\hat{Y} \equiv p(\hat{Y} \equiv Y|M). \quad (56)$$

That is, we have shown that Bayesian model testing is a special case the generalized model reliability metric in the case of exact agreement (31), i.e. $\epsilon = 0$, as can be seen by investigating (53).

Thus, a generalization of Bayesian model testing is to let the definition of agreement have a tolerance $\epsilon > 0$ such that the Bayes factor becomes,

$$K = \frac{p(A|M, D)}{p(A|M', D)} \longrightarrow K(\epsilon) = \frac{p(A|M, D, \epsilon)}{p(A|M', D, \epsilon)}, \quad (57)$$

where $p(A|M, D, \epsilon)$ is Eq. (28). This derivation suggests that the BVM can be used analogously to Bayesian model testing, except with arbitrary definitions of agreement $\epsilon \rightarrow B$, that is, we may construct the BVM factor,

$$K(B) = \frac{p(A|M, D, B)}{p(A|M', D, B)}. \quad (58)$$

Using Bayes Theorem, $p(M|A, D, B) = \frac{p(A|M, D, B)p(M|D, B)}{p(A|D, B)}$, and so we may further construct the BVM ratio,

$$R(B) = \frac{p(M|A, D, B)}{p(M'|A, D, B)} = \frac{p(A|M, D, B)p(M|D, B)}{p(A|M', D, B)p(M'|D, B)} = K(B) \frac{p(M|D, B)}{p(M'|D, B)}, \quad (59)$$

for the purpose of model testing under a general definition of agreement B , i.e., we can do model selection under any definition of agreement with the BVM ratio. The ratio $\frac{p(M|D, B)}{p(M'|D, B)}$ is the ratio prior probabilities of M and M' , which again, if there is no reason to suspect that one model is a priori more probable than another, one may let $\frac{p(M|D, B)}{p(M'|D, B)} = 1$, and then $R(B) \rightarrow K(B)$. With the BVM ratio one could in principle compare data and models with different data types to perform model testing or selection.