

AN OPTIMIZATION FRAMEWORK FOR HYBRID FIRST-PRINCIPLES DATA-DRIVEN MODELING

REMI LAM*, LIOR HORESH†, HAIM AVRON‡, AND KAREN WILLCOX§

Abstract. Mathematical models are used extensively for diverse tasks including analysis, optimization, and decision making. Frequently, those models are principled but imperfect representations of reality. This is either due to incomplete physical description of the underlying phenomenon (simplified governing equations, defective boundary conditions, etc.), or due to numerical approximations (discretization, linearization, round-off error, etc.). Model misspecification can lead to erroneous model predictions, and respectively suboptimal decisions associated with the intended end-goal task. To mitigate this effect, one can amend the available model using limited data produced by experiments or higher fidelity models. A large body of research has focused on estimating explicit model parameters. This work takes a different perspective and targets the construction of a correction model operator with implicit attributes. We investigate the case where the end-goal is inversion and illustrate how appropriate choices of properties imposed upon the correction and corrected operator lead to improved end-goal insights.

1. Introduction. Mathematical models play a central role in numerous fields: description, analysis, optimization, inversion, decision making, etc. Development of a model requires proper balance between multiple objectives that influence the model realism and its complexity. One desirable property is to achieve high fidelity: an accurate representation of the phenomenon of interest. Another objective is to develop a model that is tractable to solve either analytically or numerically. These competing requirements typically lead to a trade-off between fidelity and complexity of the model. For example, one might choose to use a simplified set of governing equations, sacrificing fidelity in favor of reduced computational complexity. As a consequence, models often form imperfect representations of the true physical process, and thus, also impact and influence the end-goal task. In this paper, we refer to such imperfect models as *misspecified models*. The inadequacies associated with such models are often disregarded, or over-simplistically treated as noise.

For the sake of clarity, we define the operators and their respective mappings. We describe the process or phenomenon under investigation by T , which is characterized by a set of model parameters m , control terms q , and experimental design θ . The potential set of state parameters u_T is given by the following equation:

$$(1.1) \quad T(u_T, q, m; \theta) = 0.$$

For simplicity, we assume that such u_T exists everywhere and is unique, and denote with P_θ the observation operator, which links a state u_T to the observations d , and is defined by:

$$(1.2) \quad d = P_\theta(u_T) + \epsilon(\theta).$$

In the above, we also incorporate a noise term, ϵ , to account for intrinsically stochastic error, which is independent of the process.

Rarely, a mathematical model describes conclusively the underlying phenomenon or process, T ; thus, in practice, often one resorts to use of an imperfect model M

*Massachusetts Institute of Technology, Cambridge, Massachusetts. (rlam@mit.edu).

†IBM Research, Yorktown Heights, New-York. (lhoresh@us.ibm.com).

‡Tel Aviv University, Tel Aviv, Israel. (haimav@post.tau.ac.il).

§Massachusetts Institute of Technology, Cambridge, Massachusetts. (kwillcox@mit.edu).

to describe the behavior of the system under investigation. Since the model M is misspecified, we shall denote by R the residual due to model discrepancy:

$$(1.3) \quad \underbrace{R(u, q, m; \theta)}_{\text{model discrepancy}} = \underbrace{T(u, q, m; \theta)}_{\text{fully specified model}} - \underbrace{M(u, q, m; \theta)}_{\text{misspecified model}}.$$

The use of a misspecified model M , instead of of the fully specified one, T , leads to an inaccurate state parameter u_M (instead of u_T) satisfying:

$$(1.4) \quad M(u_M, q, m; \theta) = 0.$$

Revisiting the observation d , this impairs the predicted data integrity through introduction of the observation error term ξ :

$$(1.5) \quad d = \underbrace{P_\theta(u_M)}_{\text{misspecified prediction}} + \underbrace{\xi(u_M, u_T, q, m; \theta) + \epsilon(\theta)}_{\zeta(u_M, u_T, q, m; \theta)}.$$

In conventional settings, both ξ and ϵ , aggregated above as ζ , are commonly regarded as an apparent “noise”. Yet, as the term noise in fact hosts various factors of different characteristics, we ought to be more careful with the terminology at this point. We therefore decouple the “noise” term, ζ , into two components: the first, $\xi(u_M, u_T, q, m; \theta)$ depends upon the parameters m , controls q and experimental design θ , whereas the latter, $\epsilon(\theta)$ depends solely upon the experimental design.

Given this setup, it is possible to decouple the effect of instrumentation noise and intrinsic stochastic effects, from misspecification of the model. This separation may seem artificial, as in principle, instrumentation error can be also included in a more comprehensive model that includes the measurement apparatus in addition to the system under investigation. Yet, in the context of this study, such separation already provides another level of granularity overlooked in the traditional setup.

Following the above notational exposition, it is now possible to distinguish between two complementary challenges:

- Model *reduction* - aims at reducing a model complexity, potentially at the expense of a decreased fidelity: $|\zeta_r(u_M, u_T, q, m; \theta)| \geq |\zeta_i(u_M, u_T, q, m; \theta)|$.
- Model *improvement* - aims at increasing a model fidelity $|\zeta_i(u_M, u_T, q, m; \theta)| \leq |\zeta_r(u_M, u_T, q, m; \theta)|$, potentially at the expense of increased complexity.

While the former has been investigated extensively (see e.g., [1, 5, 25]), little attention has been given to the development of consistent, derivable frameworks for the latter challenge.

The goal of this study is to propose a framework for inference of a *correction* model, C , for the misspecified term, M , such that the *corrected*¹ model $M(u, q, m; \theta) + C(u, q, m; \theta)$ performs better for a desired end-goal task, in a sense yet to be defined. Thus, paraphrasing George Box’s statement “All models are wrong, but some are useful”, our goal is to proactively make models more useful, through learning an approximated correction model for the misspecified elements. The framework is particularly suited for situations where a first-principles model provides an inadequate approximation of the underlying phenomenon or process. Yet, as not all governing aspects of the phenomenon are fully realized in the model (either due to lack of knowledge, understanding or computational limitations), a hybrid first-principles

¹Note that we distinguish between the correction model and the corrected model, as different requirements may be imposed upon each of them.

Attribute	Data-driven	First-Principles
Adaptability & deployability	+ Generic frameworks, easily adapted to new problems	- Requires complex derivation, specific to the application
Domain expertise reliance	+ Provide useful results with little domain knowledge	- Depends heavily upon domain expertise
Data availability reliance	- Usually requires Big Data to train upon	+ Usually can be derived from Small Data
Fidelity & robustness	- Limited generality associated with the training set span and complexity	+ Universal links of complex relations
Interpretability	- Limited due to functional form rigidity	+ Physically meaningful and interpretable link between parameters

Table 1: Common characteristics of first-principles modeling and data-driven approaches.

data-driven approach, in which the cardinal missing elements of the model are learned from data, can enhance the fidelity and utility of the model.

In Table 1, some of the common characteristics of first-principles modeling are compared against those of data-driven approaches. Rather than choosing between the two, the approach taken in this study is to combine the strengths of the two approaches together and thereby attain models of superior attributes for a given task at hand.

This paper is organized as follows. Section 2 reviews some of the prior art in the area. In Section 3, we formally define the problem, the involved spaces, and properties of the correction and corrected models. In Section 4, for the sake of illustration, we describe a particular instance of the general framework. In Section 5, we present several simplifications of the optimization problem to keep it tractable. A numerical study is presented in Section 6, to demonstrate the utility of the framework. Finally, we close with conclusions and perspectives in Section 7.

2. Prior Art. A great body of research has been dedicated, somewhat disjointly, to either first-principles modeling [22, 21], or statistical modeling [7, 6, 8]. A few interesting instances in which model fidelity is enhanced through a combination of the two can be found in the literature. There are several types of discrepancy that reduce model fidelity. Kennedy and O’Hagan [17] define them, including model inadequacy: the “difference between the true mean value of a real world process and the code value at the true value of the inputs”. They mitigate this discrepancy using Gaussian Process models to approximate the difference between a low-fidelity model output and high-fidelity data [18]. Recent work [23, 24] has tackled the inadequacy issue in fully-developed channel flows and wall-bounded flows modeling by using a stochastic inadequacy model and calibrating the uncertain parameters with available high-fidelity data.

Other studies that leverage data to improve a model include [11], which considers input-output pairs of multiple physical models as a training set for a machine learning framework. This is used to create a blended model of superior fidelity for each original model. Zhang’s thesis [27] focuses on how the collection of data influence modeling, and develops a dynamic Design of Experiments framework. There is an interesting interplay between experimental design [9, 10] and model correction. Both may modify the observation operator, in order to better reflect how observables and model parameters are related. However, in experimental design, a common assumption is that the underlying phenomenon is understood, and the challenge is to gain more understanding of the model parameters. In contrast, model correction challenges the

fundamental assumption that the governing behavior of the system is known.

Correcting a model has also been studied in the context of convolution operators, where a notable body of research is devoted to estimation of impulse response functions based on observational data in the form of blind-deconvolution [4, 3, 19].

A stochastic optimization framework for linear model correction has been proposed in [13, 12]. The approach provides means to balance between fidelity of the corrected model and complexity of a non-parameterized correction. Since often there are also properties that one may wish to attribute to the corrected model, rather than to the correction itself, in this study, we generalize the framework to allow greater control upon the designed corrected model.

3. Mathematical Framework.

3.1. Variables and Spaces Definition. In the following, we consider the simplified situation where the model parameter m and the experimental design setup θ are fixed. In this case, the fully specified model T of Eq. 1.1 becomes:

$$(3.1) \quad T(u_T, q) = 0,$$

and the misspecified model M of Eq. 1.4 is:

$$(3.2) \quad M(u_M, q) = 0.$$

The observation model Eq. 1.2 becomes:

$$(3.3) \quad d = P(u_T) + \epsilon,$$

where ϵ represents the noise. Note that d (which could be multidimensional) depends on the control q through the state² u_T , defined by $T(u_T, q) = 0$. We consider the case where we have N_T observations. Each observation $d^{(i)}$, $i = 1, \dots, N_T$ corresponds to a particular realization of the control $q^{(i)}$, $i = 1, \dots, N_T$. The collection Q of such controls and the collection D of associated observations define a training set.

We now consider the problem of estimating a correction model C , such that the corrected model behaves like the fully specified one T for a specific end-goal task.

3.2. Properties of the Correction and Corrected Operator. Given a misspecified model and a training set from the fully specified model, our goal is to form a correction model C , such that the corrected model achieves a higher *fidelity* than the misspecified one for an end-goal task. One way to achieve this goal is to design C such that the discrepancy on the training set is minimized. When little data regarding the process of interest is available, such a strategy alone is problematic because the problem may be under-determined.

Depending on the end-goal objective, some properties, or *virtues*, of the corrected model may be desirable. Here, we define a virtue to be a mapping, \mathcal{V} , from the misspecified and correction model to a real valued scalar. Possible virtues could be computational speed (for a real-time application or scalability considerations), conditioning (for matrix inversion), stability (for control), etc. Enforcing such virtues restricts the space of admissible correction models C .

Finally, we may want to impose some *structure* upon the correction. Here, we define a structure to be a mapping, \mathcal{S} , from the correction model to a real valued scalar. This structure could be explicit: e.g., a parametric form of the correction, symmetry,

²Note that the state u_T is in general not directly observable.

tri-banded structure, invariance under some operations, etc. This structure could also be implicit: sparsity of the correction, low-rankness, existence in a low dimensional manifold, etc. Low dimensional structural preferences can be instrumental in ensuring that C is ‘small’ compared to M , as it is indeed a correction.

We formulate the problem of estimating the correction model as an optimization problem involving the three aforementioned quantities: fidelity, virtue, and structure:

$$(3.4) \quad C = \underset{C}{\operatorname{argmin}} \mathcal{G}[\mathcal{F}(M, C, D, Q)]$$

$$(3.5) \quad \text{s.t.} \quad \mathcal{V}(M, C) \leq \nu$$

$$(3.6) \quad \mathcal{S}(C) \leq \tau,$$

where \mathcal{G} represents a noise model, \mathcal{F} is the functional form that quantifies the fidelity of the corrected model with respect to the training set comprising a set of controls Q and a set of corresponding observations D , \mathcal{V} quantifies the virtue of the corrected operator, and \mathcal{S} qualifies the structure of the correction C . The values $\nu \in \mathbb{R}$ and $\tau \in \mathbb{R}$ are user-specified constants that specify the desired restrictions on the virtue metric and the structure metric, respectively. The next section gives an illustrative example for choices of these quantities. An important aspect of ongoing work is to identify particular choices for quantification of fidelity, virtue and structure that lead to a well-posed optimization problem Eq. 3.4–3.6.

4. An Example of Model Choices. While the framework described in Section 3.2 is general, we now make some model choices for the purpose of illustration.

We consider the case where the misspecified model M is characterized by an invertible linear operator. In this case, the general misspecified model Eq. 3.2 is given by:

$$(4.1) \quad M(\mathbf{u}_M, \mathbf{q}) = \mathcal{M}\mathbf{u}_M - \mathbf{q} = 0,$$

where \mathcal{M} is a square invertible matrix mapping a state vector \mathbf{u}_M to a control vector \mathbf{q} . We restrict our attention to the class of additive correction model C , characterized by linear operators \mathcal{C} ; thus, a corrected operator acting upon a state vector \mathbf{u} can be written $(\mathcal{M} + \mathcal{C})\mathbf{u}$.

The correction of a misspecified model is tailored to a given end-goal task. In this example, we choose this task to be the resolution of the inverse problem of estimating the state given the control \mathbf{q} . We are given a training set of observed data and controls $\{D, Q\}$ from the fully specified model, where D (respectively Q) is a matrix whose i^{th} column is the i^{th} observation vector $\mathbf{d}^{(i)}$ (respectively control vector $\mathbf{q}^{(i)}$) of the training set. We consider the case where the data observed is a subset of the state field, so the observation operator is represented by a linear sampling operator \mathcal{P} (for instance a matrix with entries 0 and 1) that selects the corresponding observed part of the state.

We now specify the three components of the optimization framework: fidelity, virtue, and structure. Given a misspecified model M , a correction C , and a training set D, Q , we define \mathcal{F} to be the error in the observable data: $\mathcal{F}(M, C, D, Q) = \mathcal{P}^\top[(\mathcal{M} + \mathcal{C})^{-1}Q] - D$. We also choose the noise model \mathcal{G} to be the square of the Frobenius norm: $\mathcal{G}[\cdot] = \|\cdot\|_F^2$.

We can now define the fidelity term, which we also refer to as the *inverse error*, to be:

$$(4.2) \quad \mathcal{G}[\mathcal{F}(M, C, D, Q)] = \|\mathcal{P}^\top[(\mathcal{M} + \mathcal{C})^{-1}Q] - D\|_F^2.$$

Note that computing the inverse error requires the corrected model to be invertible.

A natural virtue to impose upon the corrected model for the task of inversion is a low condition number for the corrected operator:

$$(4.3) \quad \mathcal{V}(M, C) = \kappa(\mathcal{M} + C).$$

A low condition number, κ , avoids ill-conditioned inversion and numerical issues. In addition, the condition number quantifies the spread of the eigenvalues, which affects the convergence rate of iterative methods (e.g., GMRES). Thus, when dealing with large systems, a low condition number is likely to reduce the number of solver iterations or equally, to offer greater precision for a similar computational budget.

For structure, we choose to impose a low-rank structure on the correction operator C . This corresponds to ensuring that the correction is small. Thus we have,

$$(4.4) \quad \mathcal{S}(C) = \text{rank}(C).$$

These are specific choices to demonstrate the proposed concepts of fidelity, virtue, and structure. Other choices are possible, as mentioned in Section 3.2, depending on the context of the problem.

5. Optimization Problem. For the sake of tractability, we now make additional simplifications to the optimization problem defined both by Eq. 3.4 and the choices and assumptions of Section 4 (Eq. 4.2–4.4).

First, we convert the virtue constraint to a regularization term and write the optimization as:

$$(5.1) \quad C = \underset{C}{\text{argmin}} \quad \|\mathcal{P}^\top[(\mathcal{M} + C)^{-1}Q] - D\|_F^2 + \lambda\kappa(\mathcal{M} + C)$$

$$(5.2) \quad \text{s.t.} \quad \text{rank}(C) \leq \tau,$$

where λ is a regularization parameter.

We now simplify the formulation in two ways. First, we use the square of the Frobenius norm of the inverse of the corrected operator $\|(\mathcal{M} + C)^{-1}\|_F^2$ as a proxy for the condition number $\kappa(\mathcal{M} + C)$. This is motivated by the idea that the Frobenius norm of an inverse matrix K^{-1} provides a lower bound on the smallest singular value σ_{\min} of K as $\|K^{-1}\|_F^2 \geq \sigma_{\min}^{-2}$, and thus controls indirectly the condition number $\kappa(K)$.

Second, the constraint on the rank is approximated by a constraint on the nuclear norm, which is the tightest convex relaxation of the rank operator. This leads to the following optimization problem:

$$(5.3) \quad C = \underset{C}{\text{argmin}} \quad \|\mathcal{P}^\top[(\mathcal{M} + C)^{-1}Q] - D\|_F^2 + \lambda \|(\mathcal{M} + C)^{-1}\|_F^2$$

$$(5.4) \quad \text{s.t.} \quad \|C\|_* \leq \delta,$$

where $\|\cdot\|_*$ denotes the nuclear norm and $\delta \in \mathbb{R}$ is the specified limit on the nuclear norm of C .

The gradient of the objective function can be computed in closed form using the following equations, where we define $A = \mathcal{M} + C$ for compactness:

$$(5.5) \quad \frac{\partial}{\partial C} \|\mathcal{P}^\top[A^{-1}Q] - D\|_F^2 = -2A^{-\top} \mathcal{P}[\mathcal{P}^\top[A^{-1}Q] - D]Q^\top A^{-\top}$$

$$(5.6) \quad \frac{\partial}{\partial C} \|A^{-1}\|_F^2 = -2A^{-\top} A^{-1} A^{-\top}.$$

The constraint Eq. 5.4 can be handled in various ways, such as semidefinite programming, projection and proximity, or Frank-Wolfe convex reformulation [20, 26, 2, 14].

6. Numerical Study: Wings Interaction. We proceed with a validation study with a simple illustrative problem in which the fully specified model \mathcal{T} is known.

6.1. Problem setup. We consider a problem where an air velocity field \mathbf{v} is induced by two wings moving at velocity \mathbf{v}_w (Fig.1). The compressibility of air is neglected, and the curl of the velocity field is assumed to be zero (no vorticity in the flow field). Under these assumptions, the velocity vector describing the flow field can be represented as the gradient of a scalar velocity potential u : $\mathbf{v} = \nabla u$. This leads to the following system of partial differential equations with non-penetration boundary conditions:

$$(6.1) \quad \nabla^2 u = 0,$$

$$(6.2) \quad \frac{\partial u}{\partial \mathbf{n}} = \nabla u \cdot \mathbf{n} = \mathbf{v}_w \cdot \mathbf{n},$$

where \mathbf{n} is the outward normal vector at the surface of a wing.




Figure 1: Value of u on the two interacting wings computed with \mathcal{T} .




Figure 2: Absolute error in u on the wing of interest computed with \mathcal{M} (only one wing).

The goal of the inverse problem is to estimate the potential u on one of the two wings for a given wing velocity \mathbf{v}_w (angle and amplitude). Accurately estimating the potential flow is important because it can be used to estimate the velocity field \mathbf{v} and predict quantities such as drag. Using a panel method [15], it is possible to compute u at the surface of the wing of interest (i.e., without computing u on a mesh as in finite volume methods). This leads to the following formulation for the fully specified model T :

$$(6.3) \quad T(\mathbf{u}_T, \mathbf{q}) = \mathcal{T}\mathbf{u}_T - \mathbf{q} = 0,$$

where \mathbf{u}_T and \mathbf{q} are vectors such that the j^{th} entry $u_{T,j} = u(\mathbf{r}_j)$ and $q_j = \mathbf{v}_w \cdot \mathbf{n}(\mathbf{r}_j)$, with \mathbf{r}_j the centroid location of the j^{th} panel and $\mathbf{n}(\mathbf{r}_j)$ the corresponding outward normal vector. \mathcal{T} is a $p \times p$ dense matrix induced by a kernel, \mathcal{K} , related to the governing equations, where p is the total number of panels used: $\mathcal{T}_{ij} = \mathcal{K}(\mathbf{r}_i, \mathbf{r}_j)$.

In this particular example, $\mathcal{T} \in \mathbb{R}^{396 \times 396}$ is the operator corresponding to the two interacting wings ($p = 396$). The misspecified model $M(\mathbf{u}_M, \mathbf{q}) = \mathcal{M}\mathbf{u}_M - \mathbf{q} = 0$, with $\mathcal{M} \in \mathbb{R}^{198 \times 198}$, is defined only with the panels on the wing of interest ($p = 198$) and is thus unable to represent the interaction between the two wings (Fig.2). In




Figure 3: Entries of the misspecified operator \mathcal{M} .




Figure 4: Entries of the correction operator \mathcal{C} .

realistic systems, one could decide to neglect such interactions for several reasons, such as the interaction may be believed to be weak, or the dimension reduction of the matrix used for inversion might lead to computational savings.

We use $N_T = 50$ different angles for the velocity \mathbf{v}_w (between -45 and 45 degrees) to generate 50 vectors \mathbf{q} . The fully specified operator is used to compute the associated vectors \mathbf{u}_T , which are used to generate a data set. The 50 data vectors \mathbf{d} are noiseless observations of the state restricted to the panels corresponding to the wing of interest. Those 50 pairs of truncated vectors are used to define a training set $\{D, Q\}$. The test set $\{D_{test}, Q_{test}\}$ is composed of 100 pairs similarly computed. The test set is used to assess the performance of the corrected model, and is not used to construct the correction operator \mathcal{C} . We apply a modified version of the Frank-Wolfe algorithm [16, 13] to solve the constrained optimization problem Eq. 5.3.

The model correction problem is solved for several values of the regularization parameter λ to infer the additive low-rank, linear, correction operator \mathcal{C} . Entries of the misspecified operator \mathcal{M} , as well as the correction operator \mathcal{C} for $\lambda = 10^7$ are shown in Fig. 3–4. The misspecified operator exhibits three dominant diagonal structures. We notice that the correction \mathcal{C} is characterized by three diagonal structures similarly located but of different values compared to \mathcal{M} . This structure arises naturally (i.e. without parameterization) from the formulation.

6.2. Proxy for the Condition Number. We first examine the condition number of the corrected operator and compare it to the proposed proxy. As expected, the condition number κ of the corrected operator (Fig. 5) behaves qualitatively similarly to its inverse Frobenius norm proxy (Fig. 6) for different values of λ . In particular, for $\lambda \approx 1$, both the condition number and the proxy decrease as λ increases, for regularization parameters larger than 10^5 both quantities are little affected.

6.3. Error Analysis. We assess the performance of the proposed method on the test set $\{D_{test}, Q_{test}\}$. We show the inverse error (Fig. 7) obtained with corrected operators $\mathcal{M} + \mathcal{C}$ (for different λ) and compare it to results computed with \mathcal{M} .

We note that, in this case, the corrected models outperform the misspecified model for all regularization parameters λ . For low λ , the constraint on the nuclear norm (Fig. 8) is not active. Thus, increasing the regularization parameter decreases the condition number (Fig. 5), yet, has little effect on the inverse error (Fig. 7). For $\lambda \geq 10^3$, the constraint becomes active and further reduction on the condition number is observed, with a decrease in inverse error. For $\lambda \geq 3 \times 10^4$, the condition number decreases further and little emphasis is given to reduction of the inverse error on the training set, which explains the slight increase of the inverse error on the test set.




Figure 5: Condition number of the corrected operator as a function of the regularization parameter λ (normalized by $\kappa(\mathcal{M})$).




Figure 6: Proxy for the condition number of a corrected operator as a function of the regularization parameter λ .




Figure 7: Inverse error on the test set as a function of the regularization parameter λ .




Figure 8: Nuclear norm of the correction operator as a function of the regularization parameter λ (normalized by δ).

Empirically, penalizing the inverse error on a training set to construct a correction operator increases performance on a test set. Enforcing a low condition number can further reduce the inverse error for a large range of λ . We note that even for $\lambda \geq 3 \times 10^4$, the inverse error is lower than what would have been obtained by not enforcing the condition number virtue (the limiting case where $\lambda \rightarrow 0$).

7. Future Perspectives. Our main contribution is an optimization framework in which models seeded from first-principles can be improved and augmented via a data-driven approach. This construction is cast as an optimization problem with three components: a fidelity term to reduce discrepancy with the data, a structure term that characterizes the form of the correction (either implicitly or explicitly), and a virtue that tailors the corrected model to a given end-goal task. This framework is quite general and choosing a specific manifestation for these three components is problem dependent.

For the purpose of illustration, we demonstrated our approach using an inversion problem with a low-rank structure for the correction and enforcing a low condition number as virtue. We have solved this particular form of the framework using the Frank-Wolfe algorithm and closed-form gradients. This was used to improve a mis-specified airfoil panel code lacking interaction between two components. Using a training set generated by a fully specified model that modeled the two interacting wings, the correction was able to adjust the model to account for that interaction.

Ongoing extensions to this work include investigating other forms of structure. For example, a sparse correction could be adapted to correct PDE operators while

maintaining their sparsity. Another extension of the framework is to nonlinear models, which would significantly extend the applicability of the proposed approach to real-world applications.

REFERENCES

- [1] A.C. Antoulas, *Approximation of large-scale dynamical systems*, SIAM, Philadelphia, PA, 2005.
- [2] H. Avron, S. Kale, S. Kasiviswanathan, and V. Sindhvani, *Efficient and practical stochastic subgradient descent for nuclear norm regularization*, arXiv preprint arXiv:1206.6384 (2012).
- [3] G.R. Ayers and J.C. Dainty, *Iterative blind deconvolution method and its applications*, Optics Letters **13** (1988), no. 7, 547–549.
- [4] A.J. Bell and T.J. Sejnowski, *An information-maximization approach to blind separation and blind deconvolution*, Neural Computation **7** (1995), no. 6, 1129–1159.
- [5] P. Benner, S. Gugercin, and K. Willcox, *A survey of projection-based model reduction methods for parametric dynamical systems*, SIAM Review **57** (2015), no. 4, 483–531.
- [6] L. Breiman et al., *Statistical modeling: The two cultures (with comments and a rejoinder by the author)*, Statistical Science **16** (2001), no. 3, 199–231.
- [7] S. Coles, J. Bawa, L. Trenner, and P. Dorazio, *An introduction to statistical modeling of extreme values*, vol. 208, Springer, 2001.
- [8] D.A. Freedman, *Statistical models: Theory and practice*, Cambridge University Press, 2009.
- [9] E. Haber, L. Horesh, and L. Tenorio, *Numerical methods for experimental design of large-scale linear ill-posed inverse problems*, Inverse Problems **24** (2008), no. 5, 55012–55028.
- [10] ———, *Numerical methods for the design of large-scale nonlinear discrete ill-posed inverse problems*, Inverse Problems **26** (2010), no. 2, 25002–25015.
- [11] H.F. Hamann, Y. Hwang, T.G. Van Kessel, I.K. Khabibrakhmanov, S. Lu, and R. Muralidhar, *Multi-model blending*, 2015, US Patent 20,150,347,922.
- [12] N. Hao, L. Horesh, and M.E. Kilmer, *Model correction using a nuclear norm constraint*, The Householder Symposium XIX on Numerical Linear Algebra, Domain Sol Cress, Spa, Belgium, 2014, pp. 131 – 132.
- [13] ———, *Nuclear norm optimization and its application to observation model specification*, Compressed Sensing & Sparse Filtering, Springer, 2014, pp. 95–122.
- [14] E. Hazan, *Sparse approximate solutions to semidefinite programs*, LATIN 2008: Theoretical Informatics, Springer Berlin Heidelberg, 2008, pp. 306–316.
- [15] J.L. Hess, *Panel methods in computational fluid dynamics*, Annual Review of Fluid Mechanics **22** (1990), no. 1, 255–274.
- [16] M. Jaggi, *Revisiting Frank-Wolfe: Projection-free sparse convex optimization*, Proceedings of the 30th International Conference on Machine Learning (ICML-13), 2013, pp. 427–435.
- [17] M. C. Kennedy and A. O’Hagan, *Bayesian calibration of computer models*, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **63** (2001), no. 3, 425–464.
- [18] M.C. Kennedy and A. O’Hagan, *Predicting the output from a complex computer code when fast approximations are available*, Biometrika **87** (2000), no. 1, 1–13.
- [19] A. Levin, Y. Weiss, F. Durand, and W.T. Freeman, *Understanding and evaluating blind deconvolution algorithms*, Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on, IEEE, 2009, pp. 1964–1971.
- [20] Z. Liu and L. Vandenberghe, *Interior-point method for nuclear norm approximation with application to system identification*, SIAM Journal on Matrix Analysis and Applications **31** (2009), no. 3, 1235–1256.
- [21] J.C. Maxwell, *A treatise on electricity and magnetism*, vol. 1, Clarendon Press, 1881.
- [22] I. Newton, D. Bernoulli, C. MacLaurin, and L. Euler, *Philosophiae naturalis principia mathematica*, vol. 1, Excudit G. Brookman; Impensis T.T. et J. Tegg, Londini, 1833.
- [23] T.A. Oliver and R.D. Moser, *Bayesian uncertainty quantification applied to RANS turbulence models*, Journal of Physics: Conference Series, vol. 318, IOP Publishing, 2011, p. 042032.
- [24] T.A. Oliver, B. Reuter, and R.D. Moser, *Formulation and calibration of a stochastic model form error representation for RANS*, Bulletin of the American Physical Society **59** (2014).
- [25] A.T. Patera and G. Rozza, *Reduced basis approximation and a posteriori error estimation for parametrized partial differential equations*, Version 1.0, Copyright MIT 2006.
- [26] K. Toh and S. Yun, *An accelerated proximal gradient algorithm for nuclear norm regularized linear least squares problems*, Pacific Journal of Optimization **6** (2010), no. 615–640, 15.
- [27] Y. Zhang, *Improved methods in statistical and first principles modeling for batch process control and monitoring*, ProQuest, 2008.