

Inference of neutrino flavor evolution through data assimilation and neural differential equations

Ermal Rrapaj,^{1,2,*} Amol V. Patwardhan,^{1,3,4,†} Eve Armstrong,^{5,6,‡} and George M. Fuller^{7,§}

¹*Department of Physics, University of California, Berkeley, CA 94720, USA*

²*School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA*

³*Institute for Nuclear Theory, University of Washington, Seattle, WA 98115, USA*

⁴*SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA, 94025*

⁵*Department of Physics, New York Institute of Technology, New York, NY 10023, USA*

⁶*Department of Astrophysics, American Museum of Natural History, New York, NY 10024, USA*

⁷*Department of Physics, University of California, San Diego, La Jolla, CA 92093-0319, USA*

(Dated: April 29, 2021)

The evolution of neutrino flavor in dense environments such as core-collapse supernovae and binary compact object mergers constitutes an important and unsolved problem. Its solution has potential implications for the dynamics and heavy-element nucleosynthesis in these environments. In this paper, we build upon recent work to explore inference-based techniques for the estimation of model parameters and neutrino flavor evolution histories. We combine data assimilation, ordinary differential equation solvers, and neural networks to craft an inference approach tailored for non-linear dynamical systems. Using this architecture, and a simple two-neutrino-beam, two-flavor model, we compare the performances of nine different optimization algorithms and expand upon previous assessments of the efficacy of inference for tackling problems in flavor evolution. We find that employing this new architecture, together with evolutionary optimization algorithms, accurately captures flavor histories in the small-scale model and allows us to quickly explore both model parameters and initial flavor content. In future work we plan to extend these inference techniques to large numbers of neutrinos.

I. INTRODUCTION

Core-collapse supernovae and binary compact object mergers are extreme physical environments with the potential to serve as valuable laboratories at the intersection of particle theory, dense matter physics, and high-energy astrophysics. Many of the important physical phenomena in these environments, such as shock propagation, bulk matter outflows, and the synthesis of heavy-elements are driven in part by interactions between nuclear matter and the accompanying prodigious flux of emitted neutrinos [1–4].

In these situations, the flavor evolution of the neutrinos is a complicated, nonlinear problem, wherein the flavor histories of neutrinos with different energies and trajectories are coupled to one another. This has been shown to lead to various collective flavor oscillation phenomena [5–18]. In particular, in the last few years, it has been demonstrated that relaxing certain assumptions regarding spatial and temporal symmetries in the neutrino flavor field can lead to flavor instabilities not previously identified, and which have not been well studied [19–45] (see also the recent review in Ref. [46] and references therein).

Since the flavor evolution of neutrinos is so inextricably linked to the transport of energy and lepton number

in these environments, it is important to identify the initial conditions and physical regimes under which these noted flavor-field instabilities and collective phenomena can manifest themselves. Meanwhile, the next generation of terrestrial detectors such as DUNE [47] and Hyper-Kamiokande [48] could potentially provide a detection of a large number ($\sim \mathcal{O}(10^3\text{--}10^4)$) of neutrinos from Galactic core-collapse supernovae events. Thus it is pertinent to ask what such a detection could reveal about neutrino properties, as well as the physics of the supernova environment itself.

The last decades have seen the rapid development of machine learning (ML) algorithms that utilize “big data” to solve difficult problems such as image recognition [49–51] and natural language processing [52, 53]. The training of most ML algorithms requires large amounts of data, in part because initial conditions are often not assumed to be well known. Not surprisingly, scientific fields that typically produce large data sets have leveraged these technological advances [54, 55]. Other domains of science and engineering, however, are characterized by sparse data—sparsity that precludes the application of such learning algorithms to problems in these fields. Instead, these fields have long established research traditions which have led to the development of predictive models. The study of neutrino physics falls into this latter category.

At first glance, then, there appears to be a dichotomy in approach: data-driven machine learning on one hand, and theoretical models on the other. If no information about initial conditions—or for that matter any model knowledge—is available but large amounts of data are,

* ermalrrapaj@gmail.com

† apatward@slac.stanford.edu

‡ evearmstrong.physics@gmail.com

§ gfuller@ucsd.edu

the first approach seems very reasonable. When dealing with sparse data, however, knowledge of the model and initial conditions is invaluable. In the physical sciences and engineering fields characterized by sparse data and highly developed theoretical frameworks, functions chosen to represent any unknown parts of a model are typically dictated by the researchers’ experience and intuition. This educated guesswork might not account for all functional forms that may appropriately represent the data. This situation calls for a less biased—and more general—means of model estimation.

Recently, in calculations of neutrino flavor transformation, there have been attempts [56, 57] to combine the best from both approaches noted above, by utilizing an inference procedure known as statistical data assimilation (SDA). In an SDA procedure, real or simulated data are considered together with assumed theoretical constraints, to complete a model with one or more unknown parameters. “Completing the model” here refers to the task of simultaneously determining the evolution histories of state variables, as well as the values of unknown parameters, subject to data and physical constraints. In [56, 57] the authors used inference to predict the flavor transformation histories of two mono-energetic neutrino beams coherently interacting with each other and with a matter background potential. Simulated “measurements” of neutrino flavor at the endpoint of the evolution, as well as theoretical constraints at other points along the trajectory, were utilized to compute the evolution histories by approximating the value of the neutrino beam flavor on a grid of uniformly-spaced radial locations. Several distinct functional forms for the unknown parameters were tested. Also in [57] the effects of incrementally adding theoretical constraints to improve the inference efficacy were analyzed.

This work explores new directions in applying SDA for constraining neutrino flavor evolution histories inside core-collapse supernovae. There are two key differences with respect to previous work. Firstly, we eliminate uniformity in grid discretization. Specifically, we include an adaptive step differential equation solver in the neural architecture [58], and we employ a deep neural network that replaces any guesswork, on the part of the researcher, of the functional forms of unknown parameters. The chosen neural architecture is a universal function approximator [59], which removes the need for guessing functional forms of unknown parameters. The solver precludes any possible model error associated with grid discretization. Secondly, we perform the calculations using nine distinct optimization algorithms, for a comparative analysis of performance.

Here we study a two-neutrino-beam, two-flavor system, with one beam initially electron flavor and the other x flavor. Given the initial flavor configuration, this toy problem can be evolved from the supernova “source” to the “detector” without significant computational difficulties [60]. Our focus, however, is not the forward evolution, but rather the estimation of model parameters re-

sponsible for a given detection.

Our results show that inference, even given such an apparently-simple setup, is an intricate task. We find that, for the specific small-scale model examined in this paper, evolutionary algorithms perform well in both the estimation of unknown parameters and initial conditions. However, this may change in more complicated setups and with large number of unknown parameters and a hybrid approach of consecutive optimizations, for instance, evolutionary algorithms followed by gradient and hessian based methods, might be more appropriate. In what follows, we will discuss the potential advantages of adaptive-step-size solvers.

In Sec. II, we describe the algorithmic setup and specifics of the problem. Then we perform two sets of inference experiments. In the first set (Secs. III and IV), we estimate unknown parameters governing the matter potential, assuming the initial conditions on flavor to be known. In the second set (Secs. V and VI), we instead estimate the initial conditions, assuming the matter potential to be known. In Sec. VII we make further considerations regarding the comparison of different models and estimation of uncertainties. We present conclusions in Sec. VIII.

II. INVERSE PROBLEMS AND NEURAL NETWORK OPTIMIZATION

A. General Framework

Data assimilation is an inverse problem formulation [61]: a procedure whereby information in measurements is used to complete a model of the system from which the measurements were obtained. For our purposes, the model \mathbf{F} is written as a set of ordinary differential equations that evolve in affine parameter r as:

$$\frac{d\mathbf{u}}{dr} = \mathbf{F}(\mathbf{u}, r, \boldsymbol{\theta}) \quad (1)$$

where the vector \mathbf{u} is the observable being modeled, with initial value \mathbf{u}_0 . The affine parametrization r may be, for example, time or distance. Any unknown parameter that influences the forward problem (differential equation) is denoted by $\boldsymbol{\theta}$. An observation \mathbf{u}_D is made at a detector location R and one seeks to estimate $\boldsymbol{\theta}$ that best fits this observation. This is achieved through minimization of the cost function:

$$\text{Cost}(\boldsymbol{\theta}) = [\mathbf{u}_D - \mathbf{u}_\theta(R)]^T \mathbf{W} [\mathbf{u}_D - \mathbf{u}_\theta(R)] \quad (2)$$

where $\mathbf{u}_\theta(R)$ is the prediction from \mathbf{F} . In practice, we may only observe a subset of the components of the vector \mathbf{u}_D , while through \mathbf{F} all components are evolved in order to predict the final values at the detector. The sparse matrix \mathbf{W} is introduced to select these components when computing the cost function.

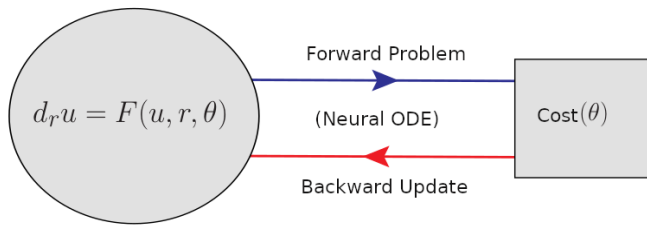


FIG. 1. Data assimilation as a neural network architecture. The forward problem is represented by a set of first-order ordinary differential equations, $F(u, r, \theta)$, that depend on the unknown model parameters θ . $\text{Cost}(\theta)$ is the objective function whose minimum denotes the optimal solution of the inverse problem. Blue arrows denote the forward process of prediction, and red arrows denote the backward pass for error correction.

In [56, 57], the cost function was comprised of three parts: model error, measurement error, and physical constraint terms. The model error, in addition to permitting uncertainty in the model parameter estimates, included terms related to the uniform discretization of the domain and finite difference approximation of the derivative. In addition, as the optimization algorithm took the grid points to be independent, we had imposed co-variation of the model coordinates into the cost function as an equality constraint. Here, we explicitly include only the measurement term in the cost function, as the neural architecture ensures that the other terms are automatically satisfied; this will be explained in detail below. In [57], we also considered u_0 to be an input in the form of a measurement. Here, in the first two experiments we follow the same assumption, but without explicitly including u_0 as part of the cost function (the cost function implicitly depends on u_0 , via the dynamical equations). In the last two experiments, we assume all model parameters are known, and we instead optimize the cost function by varying the initial conditions u_0 .

In recent decades, machine learning has been used to provide solutions to ordinary differential equations [62–66]. As the focus of this work is the inverse problem, the system of differential equations is a building block of our setup. As such, our goal is not to approximate the ODE solution through a neural architecture or discretized grid, but rather to understand which parameters in the ODE definition lead to a solution that best matches observations. Hence, we use the existing vast and established literature on solving ODEs through forward integration in the vein of incorporating model knowledge with machine learning. This is the motivation for using the recently developed neural ODE [58] network for data assimilation. This network will automatically incorporate our model knowledge of the dynamical system, through Eq. (1).

We include an adaptive step solver method in the neural architecture, removing the need for domain discretization and errors induced by such discretization. Specifically, we employ the Radau method [67]. The forward-problem arrow in Fig. 1 refers to this part of the archi-

tecture. The solution $u_\theta(r)$ then satisfies all the physical constraints associated with Eq. (1). Consequently, no model and no constraint terms are needed in the cost function. As such, the cost function in Eq. (2) contains only the measurement term. We have verified that the errors associated with model and constraint terms are within numerical precision ($\lesssim 10^{-16}$). An additional feature of this setup is the reduction of the number of unknown parameters that require optimization. Obtaining adequate resolution in the previous setup required a rather large number of grid points. In this new architecture, the points in the domain are automatically chosen by the adaptive Radau method to solve the forward problem to machine precision.

At first glance, the choice to incorporate the ODE solver into the neural architecture may seem puzzling. The most general definition of a neural layer is a differentiable function with tensor input and output. This is the basis of a general differentiable programming architecture. Traditionally, neural layers are superpositions of simple primitive functions, but they need not be. Differential equation solvers naturally fit this framework, as an ODE solver has an input vector $u(r_n)$ that outputs a new vector $u(r_{n+1})$, where the points (r_n and r_{n+1}) and the separation between those points are determined adaptively. In order for the solver to be a neural layer, the output of the ODE solver must be differentiable with respect to the unknown parameters. This is achieved through automatic differentiation [68, 69] and adjoint sensitivity analysis [70, 71]. Automatic differentiation in computer science encompasses a set of techniques for converting a program into a sequence of primitive operations that have specified routines for computing derivatives. It is efficient, in that there exists a linear time cost in computing values, and it is numerically stable. Adjoint sensitivity analysis allows for the automatic differentiation of ordinary differential equations. The code for this work was written in the Julia programming language [72], and we used the DiffEqFlux package [73–75]. While the interested reader may delve deeper into these interesting topics, for the purposes of this work, it suffices to state that the current architecture allows us to compute the Jacobian and Hessian matrices of the dynamical model ($\nabla_\theta \text{Cost}(\theta)$ and $\partial_{\theta_i} \partial_{\theta_j} \text{Cost}(\theta)$).

Once the forward problem is implemented, one performs the minimization of the cost function. In Fig. 1 this step is represented by the backward update arrow. Typically, in deep learning, the minimization is performed through stochastic gradient descent (SGD) [76] (which requires the computation of the gradients mentioned). In practice, the step size of the parameter update is a hyper parameter of the training procedure, which must be tuned to achieve convergence. Many improvements on SGD have been developed through decades. For instance, AdaGrad [77] and Adam [78] are algorithms widely used in deep learning. LBFGS [79], which also approximates the Hessian of the cost function, is a commonly used alternative to SGD.

The focus of this work is global optimization (finding the optimal value in the entire region of interest), and the algorithms mentioned above work locally, informed by gradients of the cost function. Hence, we also employed a number of global algorithms in our analysis. A list of algorithms used in this work, along with their classification, is given below:

- Monte Carlo based methods (gradient free): “Simulated Annealing” (SAMIN) [80, 81] — based on Metropolis-Hastings algorithm to generate samples from a thermodynamic system,
- Evolutionary algorithms (gradient free):
 1. “Improved Stochastic Ranking Evolution Strategy” (ISRES) [82, 83] — based on a combination of a mutation rule (with a log-normal step-size update and exponential smoothing) and differential variation (update rule similar to Nelder-Mead [84]) ,
 2. “Adaptive Particle Swarm Algorithm” (APS) [85] — improve global coverage and convergence by switching between four evolutionary states: exploration, exploitation, convergence, and jumping out.
- Jacobian and Hessian based methods:
 1. “Interior Point Optimizer” (IPOPT) [86] — a primal-dual interior point method which uses line searches based on filter methods. IPOPT is designed to exploit 1st and 2nd derivative information if provided. If no Hessians are available, IPOPT will approximate them using a quasi-Newton methods, specifically a BFGS update.
 2. “Newton method with Trusted Region Hessian” (NTR) [87] — quadratic approximation of the objective function by means of the hessian with steps restricted to be within a ‘trusted’ region where the approximation is believed to be valid.
- Combination of global and local optimization:
 1. “Stochastic Global Optimization” (STOGO) [88] — systematically divide the search space (which must be bound-constrained) into smaller hyper-rectangles via a branch-and-bound technique, and searching them by a gradient-based local-search algorithm.
 2. “Multi-Level Single-Linkage” (MLSL) [89, 90] — global optimization by a sequence of local optimizations from random starting points, in conjunction with local optimizations algorithms

- BOBYQA [91] — (gradient free) bound-constrained optimization using an iteratively constructed quadratic approximation for the objective function,
- “Method of Moving Asymptotes” (MMA) [92] — local, convex and separable approximation of the objective function from the gradient,
- LBFGS [93, 94] — quasi-Newton method that approximates the Broyden-Fletcher-Goldfarb-Shanno algorithm (BFGS) [95] using a limited amount of computer memory.

These algorithms cover a wide range of methodologies. For instance, multi-level algorithms (MLSL) have been used for over three decades in optimization, and nowadays are part of many statistical programming languages. Simulated annealing is another widely used algorithm, with over four decades of applications. In addition to the familiar Newton’s method and IPOPT, we have also included evolutionary algorithms like ISRES and APS which are inspired by biological evolution [96]. Clearly, optimization is a rather fascinating and varied field. The numerical implementation can be found in NLOpt [97] and Optim [98] packages. In all experiments covered in this work, we set the maximal number of iterations for the optimization procedure to 1000.

B. Specifics of the problem

The neutrino flavor evolution problem has been explained in detail in [57]. Here we summarize the system of differential equations,

$$\mathbf{F}_i = \frac{d\mathbf{P}_i}{dr} = \left(\Delta_i \mathbf{B} + V(r)\hat{z} + \mu(r) \sum_{j \neq i} \mathbf{P}_j \right) \times \mathbf{P}_i \quad (3)$$

Here, $\Delta_i = \delta m^2/(2E_i)$ are the vacuum oscillation frequencies of neutrinos with energies E_i . The mass-squared differences in vacuum are δm^2 . The unit vector representing neutrino flavor mixing in vacuum is $\mathbf{B} = \sin(2\alpha)\hat{x} - \cos(2\alpha)\hat{z}$, where α is the mixing angle between the flavor and mass eigenstates. The functions $V(r)$ and $\mu(r)$ are the potentials arising from neutrino-matter and neutrino-neutrino interactions, respectively. The “polarization vectors” \mathbf{P}_i , which contain information about the flavor composition of the neutrinos, play the role of the state variable \mathbf{u} from Eq. (1), and the only components that are measured at the detector are the P_z of each neutrino. Anti-neutrinos are not considered at present, and will be included in future work. The generalization to include anti-neutrinos is straightforward, and simply involves allowing for negative oscillation frequencies Δ_i in Eq. (3).

In our model, we take the neutrino-neutrino potential to be,

$$\mu(r) = \frac{\mu_0}{(r + \delta_0)^4}. \quad (4)$$

This choice is consistent with how this coupling strength varies in the neutrino bulb model calculations employing the single-angle approximation. In our SDA experiments, μ_0 is taken to be a constant with a known value and $\delta_0 = 10^{-3}$ is added to avoid any numerical singularities at $r_0 = 0$. The matter potential $V(r)$ is chosen to be

$$V(r) = \frac{V_0}{(r + \delta_0)^3}. \quad (5)$$

In the first two experiments, V_0 is treated as an unknown parameter that we optimize. To generate the simulated “detector data” (\mathbf{u}_D in Eq. (2)), we use a constant value $V_0 = \tilde{V}_0$ of the matter potential coefficient, given in table I. As a thought experiment and proof of concept, we study a system of two neutrino beams, with all the parameters used for simulated data generation displayed in table I. In future work we intend to study much larger systems.

Parameter	Value	Initial polarization	Value
Δ_1	30	$P_{1,z}(r_0)$	-1.0
Δ_2	55	$P_{2,z}(r_0)$	1.0
μ_0	10.0	Final polarization	Value
\tilde{V}_0	50.0	$P_{1,z}(R)$	0.20575
α	0.15	$P_{2,z}(R)$	-0.96750
r_0	0		
R	5		

TABLE I. **Model parameters used for generating the simulated ‘detector’ data.** Δ_i are the vacuum oscillation frequencies of the neutrinos, and (μ_0, \tilde{V}_0) are the multiplicative factors governing the neutrino-neutrino coupling potential $\mu(r)$ and matter potential $V(r)$. Parameter α is the mixing angle in vacuum. Neutrino 1 is initially x flavor and neutrino 2 is initially electron flavor, as is reflected in the respective initial P_z values.

Fig. 2 displays the solution to the forward problem—represented here by the z components of the polarization vectors as functions of r , for the model parameters of table I. The two neutrino beams are initially in pure flavor states (electron and x). At some intermediate distance, there occurs a large flavor transformation, and the two neutrino beams interchange flavors. As expected, with increasing distance, both matter and neutrino potentials become less relevant and we can observe vacuum oscillations.

III. MATTER POTENTIAL COUPLING AS AN UNKNOWN CONSTANT

In this section, we assume the matter potential coupling V_0 is an unknown parameter and ask whether the

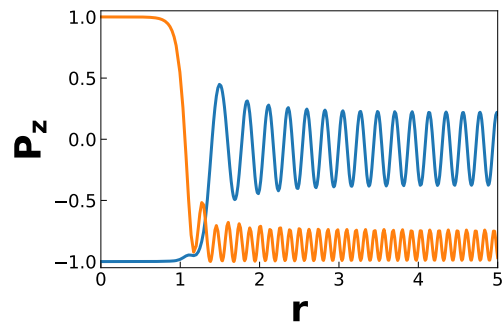


FIG. 2. The solution to the forward problem, in terms of the z components of the polarization vectors, \mathbf{P}_z , of the two neutrino beams, as functions of the affine parameter \mathbf{r} . The parameters used in the forward integration are shown in table I. One neutrino beam (orange) is initially electron flavor, and the other (blue) is initially x flavor. In this toy problem, we assume the detector to be at location $r = 5$. For solving the inverse problem, the only information available to the procedure is the polarization measured at this location.

z components of the polarization vectors of the two neutrino beams at the detector provide sufficient information to infer the true value \tilde{V}_0 provided in Table I. As this is a small system, and there is only one unknown parameter, we can plot the dependence of the cost function on the unknown parameter using repeated forward integration, as shown in Fig. 3. That is, the forward code was run several times with different parameter values V_0 , and the corresponding values of P_z at the endpoint in each case were compared with the true values (that is, with $V_0 = \tilde{V}_0$) to generate the cost function using Eq. (2). In practice, physical systems contain many more particles, and more unknowns, rendering it infeasible to create the plot analogous to Fig. 3. We are considering a ‘toy’ problem, however, as a proof of concept for the approach we propose, and for the relative ease with which we may examine figures such as those displayed in this work.

From Fig. 3 we see that at $V_0 = \tilde{V}_0$ the cost function attains its global minimal value as expected. In addition, though, there are many local minima present, which will make it difficult for any local optimization algorithm to find the correct value \tilde{V}_0 if the initial guess is in the vicinity of a different local minimum. In addition, in the parameter range with small V_0 values, the cost function changes very quickly between large and small values. To understand how gradient and Hessian based optimizers would perform, in Fig. 4 we plot the first and second derivatives of the cost as function of V_0 . As the figures show, there are large fluctuations in the derivatives of the cost function at small values of V_0 , and many stable local minima at large V_0 . On the one hand, if the initial guess is small, gradient-based optimization will change the value of the guess drastically, and on the other hand, if the initial guess is large, the changes will be minuscule. This does not bode well for gradient- (and Hessian-)

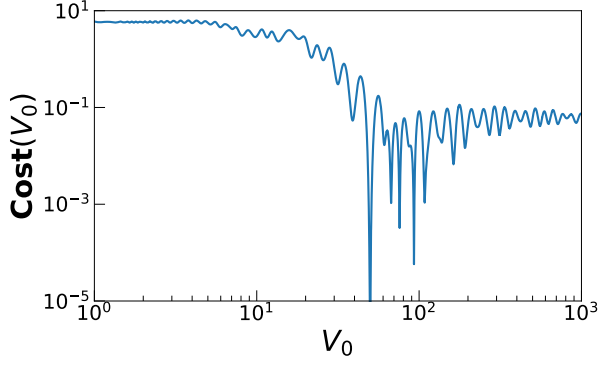


FIG. 3. Cost function dependence on the free parameter V_0 , where the simulated “detector” data was obtained through forward integration using the parameter values of table I. The global minimum is obtained at $V_0 = \tilde{V}_0$ (see table I), but there are many local minima present. Thus, even with only a single unknown parameter, this problem presents a very difficult challenge for the optimization.

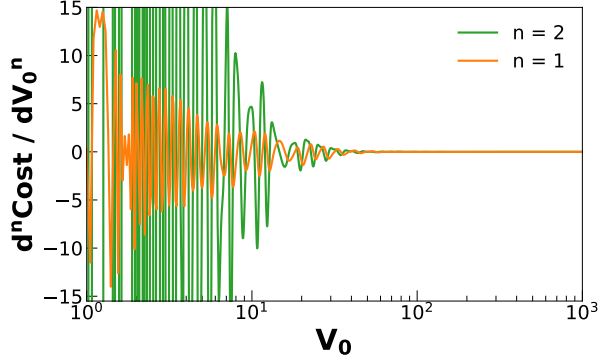
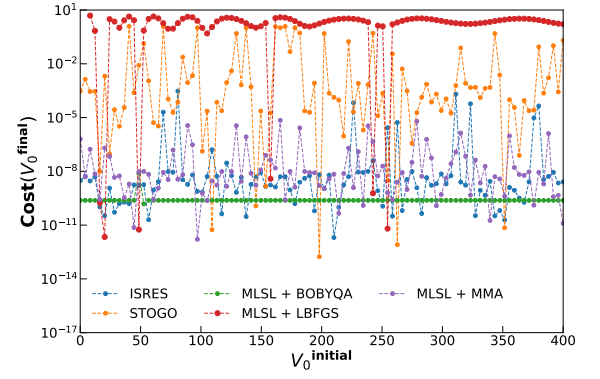


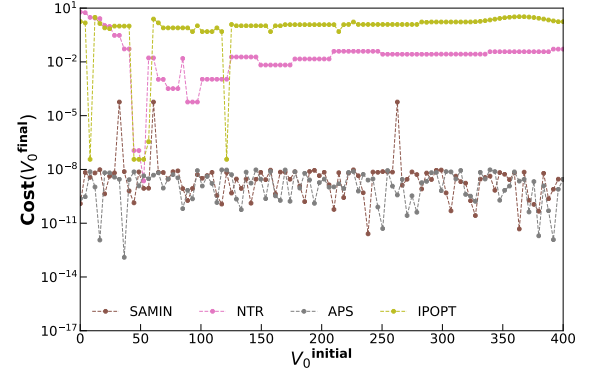
FIG. 4. Gradient (orange) and Hessian (green) of cost function with respect to the free parameter $\theta = V_0$. Ideally, one expects that these derivatives will be large at locations far from the global optimum, and that they will gradually go to zero as the minimum is approached. In our setup there is a sudden transition from large oscillations to tiny ones near $V_0 = \tilde{V}_0$. For small V_0 , gradient based methods will have very large parameter updates in either direction, and for large V_0 the updates will be very small.

based optimization. Even with just two neutrino beams, the problem is rather complicated. Hence our choice for global optimization and the wide range of optimization algorithms that we test.

In Fig. 5 we plot the cost function at the end of the iterations for each algorithm as a function of the initial guess for V_0 . We sampled uniformly 100 initial values in the range $[0, 400]$. Not surprisingly, gradient and Hessian based algorithms have a final cost value much larger than the rest. To verify that small cost indeed translates to convergence to the global minimum, we also plot the final inferred values of the unknown parameter as function of the initial guess in Fig. 6.



(a)



(b)

FIG. 5. The cost function at the final inferred values of V_0 (i.e., at the end of the optimization procedure) as function of the initial guess V_0^{initial} , for each of the nine optimization algorithm. Ideally this value should be 0, and in practice, the smaller it is the better the optimization. For visual clarity, the nine algorithms are plotted in two groups 5a and 5b. In this set of experiments, the optimization procedure is instructed to regard V_0 as a constant unknown parameter.

As can be seen from both figures, most gradient-based methods have difficulties in converging to the optimal value, while gradient-free methods perform rather well. In particular, with the Jacobian/Hessian based methods like IPOPT and NTR, the final inferred value of V_0 is positively correlated with the initial guess, suggesting that the optimization procedure simply finds a local minimum close to the initial guess. This outcome agrees with expectations laid out by Figs. 3 and 4. In addition, the combination of a global (MLSL) and local method (MMA) seems successful.

IV. MATTER POTENTIAL COUPLING AS AN UNSPECIFIED FUNCTION OF POSITION

In the calculations presented in this section, we treat the numerator in Eq. (5), V_0 , as a function of the affine parameter r . Note that the simulated “detector data” (u_D) used is identical to the data used in the previous

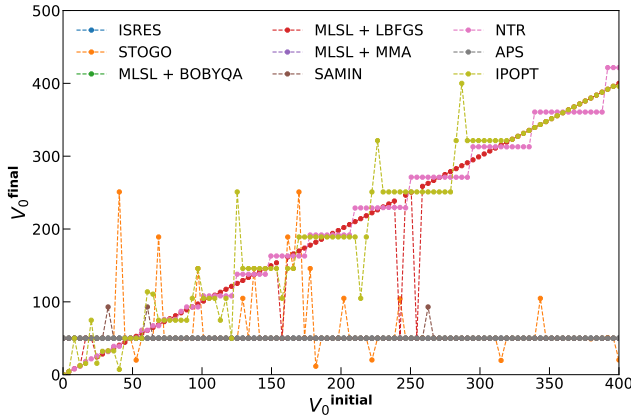


FIG. 6. The final inferred value of V_0 as function of the initial guess V_0^{initial} , for each optimization algorithms. For a converged optimization this value should be $\tilde{V}_0 = 50$, independent of the initial guess (see table I). These results are for the same set of experiments as those in Fig. 5.

section. That is, it was generated using a constant matter potential coefficient \tilde{V}_0 . Instead of selecting a specific functional form for this dependence, we represent the numerator by a two-layer feed-forward neural network, $V_0(r) = |\mathcal{N}_A(\theta, r)|$. Each layer has five neurons. The first layer has a hyperbolic tangent activation function, and the second is linear; for a total of 16 parameters denoted by θ . The depth, width, and activation functions of the neural architecture are hyper parameters. We choose these hyper parameters strictly, as our goal here is to merely explore whether a neural architecture can provide us with reasonable estimates for V_0 and the architecture chosen can represent a wide range of functions. Since only positive values are physically meaningful, the matter coupling parameter is taken to be the absolute value of the output of the neural architecture. We opted to perform global optimization in the range $\pm 10^3$ for each of the parameters. For each method we sampled uniformly 40 initial parameter sets for each optimization.

In Fig. 7 we show the maximal, minimal and average values of the cost function at the end of the each optimization procedure. APS stands out from the rest: it performs quite well for all initial guesses and provides overall small cost values. While NTR achieves a near-zero cost value for a particular initial guess, its results are quite spread and have a strong dependence on initial conditions. Generally, dependence on the initial guess is to be expected. If a guess happens to be close to the optimal result, one would expect the optimization procedure to produce a final cost value close to zero. On the other hand, if the initial guess is quite far from the optimal value, the optimization might converge to local minimal nearby. An additional complication arises from the possibility of degeneracies as the number of unknown parameters increases. “Degeneracies” here refers to the

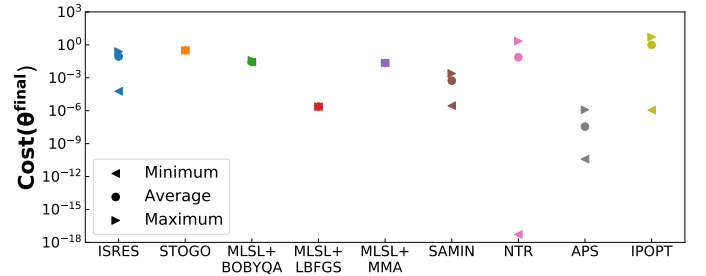


FIG. 7. The minimum, maximum and average values of the cost function at the final values of θ (i.e., at the end of the optimization procedure), for each optimization algorithm. Ideally, all these three values should be 0. The spread shows how dependent the optimization algorithms are on the initial guess of the unknown parameters. In this set of experiments, the optimization procedure treats V_0 as a generic function, encoded in terms of a two-layer feed-forward neural network, that depends on unknown parameters θ .

possibility of multiple solutions $|\mathcal{N}_A(\theta, r)|$ that yield the same values of $P_z(r)$ at the endpoint.

To illustrate the behavior behind this remark, we plot the numerator $|\mathcal{N}_A(\theta, r)|$ for each method using the parameters that produced the smallest cost value in Fig. 8. A priori, we know that a constant function of r , namely $V_0(r) = \tilde{V}_0$, is one possible optimal solution.

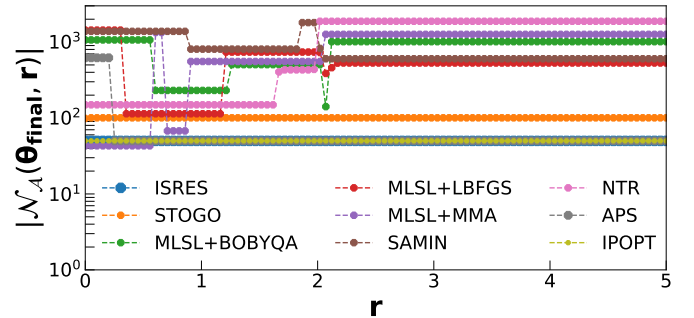


FIG. 8. Optimal matter potential coupling V_0 found by each method as function of the affine parameter r . Some algorithms find the coupling to be a constant close to the value in table I, while others show sharp matter profile changes. This shows that various matter profiles, despite being quite different, lead to very similar values of the polarization z -components detected at $r = 5$. This is for the same set of experiments described in the caption of Fig. 7.

As the plot shows, both ISRES and IPOPT converge to a constant function close to \tilde{V}_0 . APS shows a sharp transition from a region of high density to the optimal \tilde{V}_0 . A rather interesting result is displayed by NTR, where the matter density profile experiences two sharp transitions and yet the cost value is small ($\approx 10^{-18}$). Other methods that display sharp transitions are MSLSL+LBFGS and SAMIN, which perform slightly better than ISRES. In fact, many methods show sharp changes in the matter

profile. For the model chosen, the matter density profile is inversely proportional to r^3 , so the sharp changes in the numerator amount to small changes in the profile itself. In an actual core-collapse supernova environment, such sharp transitions may represent, for instance, a dense matter outflow in a lower density background, or alternatively the presence of a shock during the supernova explosion. In this manner, allowing for a variable numerator represented by the neural architecture can lead us to discovering other possible matter profiles consistent with the same detector measurement.

V. PARTIALLY UNSPECIFIED INITIAL CONDITIONS

In this section we study the influence of initial conditions on the inference of the neutrino flavor composition at the detector. As an illustration, we assume that Neutrino 1, which in the original setup was taken to be initially in the x flavor, decouples earlier and can potentially oscillate in flavor before the second neutrino is emitted. By assuming the matter and neutrino potentials of table I, and given detector measurements at $R = 5$ in Fig. 2 (i.e., the same simulated detector data as in the previous experiments), we investigate whether we can infer the initial polarization of this neutrino. The initial flavor polarization is normalized, so it can be expressed by two free parameters, the azimuthal and polar angles in flavor space; that is, $u_0^{(1)} = \{\cos(\theta_{1A})\sin(\theta_{1P}), \sin(\theta_{1A})\sin(\theta_{1P}), \cos(\theta_{1P})\}$, where $\theta_{1A} \in [0, 2\pi]$, $\theta_{1P} \in [0, \pi]$. Here, $\theta = \{\theta_{1A}, \theta_{1P}\}$ are the unknown parameters to optimize. As we assume coherent evolution, the polarization is normalized throughout the evolution.

In Fig. 9 we plot the cost function dependence on these two parameters. Given that the second neutrino is initially of electron flavor, the optimal value for the first neutrino is to be an x flavor (i.e., $P_z = -1$, or equivalently, $\theta_{1P} = \pi$), as the figure confirms. In addition, we can see that the polar angle plays a major role in determining the value of the cost function. The figure shows the cost value decrease as the polar angle changes from 0 to π , which is the optimal value. There is also minor dependence on the azimuthal angle for a fixed polar angle. Overall, this is to be expected, as for $\theta_{1P} = \pi$, the neutrino is x flavor regardless of θ_{1A} .

We have performed 100 optimization experiments with different initial guesses of θ , from a uniform grid of initial values within the allowed range. The statistics for the final values are summarized in table II. APS and MLSL + BOBQYA reach the optimal value of θ_{1P} for most of the initial guesses, as shown by the tiny variances in the table. Many other algorithms (ISRES, MLSL + MMA, SAMIN, IPOPT) converge quite close to the optimal value and have small variances. All methods converge to large values of the azimuthal angle. To understand this behavior, we computed the gradients of the

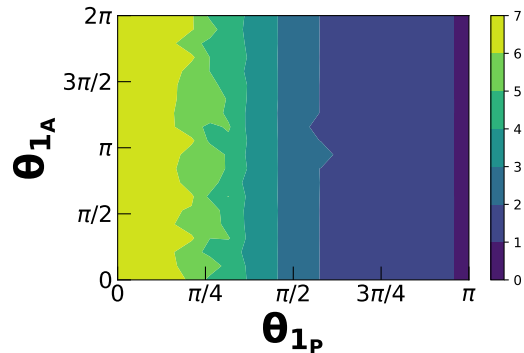


FIG. 9. Dependence of the cost function, $\text{Cost}(\theta_{1P}, \theta_{1A})$, on the initial orientation in flavor space of the polarization vector of the first neutrino. There is a relatively strong dependence on the polar angle, θ_{1P} , and weak dependence on the azimuthal angle, θ_{1A} . This is an indication of degeneracy in parameter space. Here, the matter potential was taken to be a known constant \tilde{V}_0 (see table I), and only the initial angles θ_{1P} and θ_{1A} were varied. The simulated “detector” data was generated by using $P_{1,z} = -1$ as the starting point of the forward integration.

Method	$(\mu_{\theta_P}, \mu_{\theta_A}) / \pi$	$(\sigma_{\theta_P}, \sigma_{\theta_A}) / \pi$
ISRES	(0.996, 1.05)	(0.001, 0.57)
STOGO	(0.87, 1.52)	(≈ 0 , ≈ 0)
MLSL + BOBQYA	(1.0, 1.4)	(≈ 0 , ≈ 0)
MLSL + LBFGS	(0.484, 1.0)	(0.31, 0.64)
MLSL + MMA	(0.997, 1.011)	(0.00013, ≈ 0)
SAMIN	(0.997, 0.996)	(0.002, 0.615)
NTR	(0.47, 0.98)	(0.31, 0.65)
APS	(1.0, 1.0)	(≈ 0 , 0.9)
IPOPT	(0.9, 1.1)	(0.2, 0.4)
Optimal Values	(1, -)	(0, -)

TABLE II. Sample average and standard deviation of the inferred angles in flavor space, for the initial polarization of the first neutrino at the end of the each optimization procedure. The optimal values shown at the end of the table; there is no preferred azimuthal angle. In this set of experiments, the matter coupling parameter V_0 is taken to be a known constant, and the initial polar and azimuthal angles, θ_{1P} and θ_{1A} , of the polarization vector of the first neutrino are regarded as being unknown parameters that the optimization procedure is tasked with inferring.

cost function and found one major gradient flow toward $\theta_{1P} = \pi$, as expected. We also found a rather small flow toward $\theta_{1A} = 2\pi$. But interestingly, the average inferred values of the azimuthal angle seem to be clustering around π rather than 2π . This is an unexpected result, and a priori hard to guess, as one would need to solve to the flavor evolution equations for all initial conditions to notice this secondary flow.

In Fig. 10 we display the maximal, minimal and average cost value obtained from each method. APS and MLSL + BOBQYA result in small cost values for all initial guesses, in agreement with table II. Most methods do

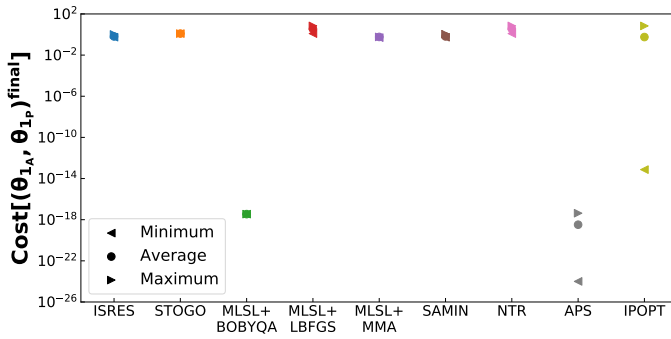


FIG. 10. The minimum, maximum and average values of the cost function at the final values of θ (i.e., at the end of the optimization procedure), for each optimization algorithm. Ideally, all these three values should be 0. For each algorithm, the spread of cost function values shows how dependent the optimization algorithms are on the initial guess of the unknown parameters. These results are for the set of experiments described in the caption of table II, where the unknown parameters to be optimized are $\theta = \{\theta_{1P}, \theta_{1A}\}$.

not show any spread in the final values of the cost function, apart from APS and IPOPT. In these two cases, for some initial conditions, the methods achieve very low cost function values.

VI. COMPLETELY UNSPECIFIED INITIAL CONDITIONS

In this section, we make no assumptions about the initial polarizations. Instead, given the detector data generated by the parameters in table I, we optimize the cost function for the azimuthal and polar angles in flavor space for both neutrinos. We pick 5 uniformly-spaced values for each of the 4 angles for a total of 625 experiments for each optimization method. These are initial guesses of the angles in flavor space.

As Fig. 11 shows, most methods converge to sub-optimal solutions. There is convergence for initial conditions close to optimal values, but this does not happen for initial guesses farther away. APS is the only method that performs well for all initial guesses. On the hand, STOGO does not provide a small cost value even for initial guesses close to the optimal one.

As an additional check, in table III we summarize the statistics for the final values obtained from each method for each of the two neutrinos. The sample standard deviation shown is an additional indication of the dependence on the initial guess. Ideally, this deviation should be zero as the methods should converge to the optimal value regardless of the initial guess for the unknown parameter, but this is not the case in practice.

As the tables show, most methods tend toward a large polar angle for the first neutrino and a small value for the second one. In other words, most methods expect the first neutrino to be mostly x flavor and the second to be

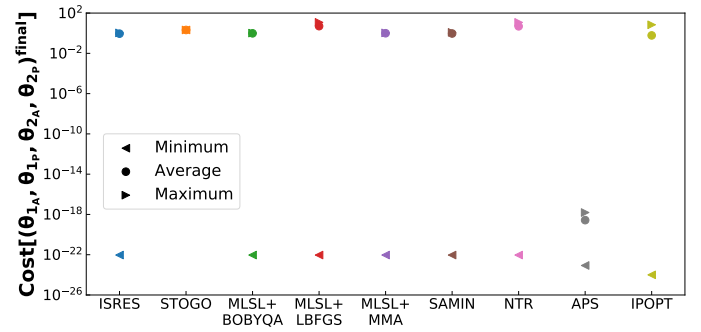


FIG. 11. The minimum, maximum and average values of the cost function at the final values of θ (i.e., at the end of the optimization procedure), for each optimization algorithm. Ideally, all these three values should be 0. For each algorithm, the spread of cost function values shows how dependent the optimization algorithms are on the initial guess of the unknown parameters. These results are for the set of experiments where the matter coupling parameter is taken to be a known constant, and the unknown parameters to be optimized are the polar and azimuthal angles of both the initial neutrino polarization vectors, i.e., $\theta = \{\theta_{1P}, \theta_{1A}, \theta_{2P}, \theta_{2A}\}$.

mostly electron flavor. This result indicates that, for the two neutrino beam system, the final polarization values can provide information on initial conditions, under the assumption that we know the matter density profile. In addition, we have identified APS as a method that works quite well in understanding the initial flavor composition of the system.

VII. FURTHER CONSIDERATIONS

In this section we ponder on further complications one may encounter when performing inference on the detector data. As the APS method was the most successful one in the various experimental setups, the additional experiments and computations performed here will focus on this method for illustration.

A. Using inference to select theoretical models consistent with detector data

In section IV we consider the matter potential coupling to be a generic function of radius. This is different from the setup used to generate the detector data which has a constant coupling \bar{V}_0 . The constant coupling is a special case for this section, and indeed, some of the optimization methods converge to it. When the detector data will be available to the community, one can not know for certain whether the theoretical model employed captures all the relevant aspects of neutrino flavor oscillations emitted from compact objects, and more probably, there will be more than one model to consider. To simulate this

Method	$(\mu_{\theta_P^{(1)}}, \mu_{\theta_A^{(1)}}) / \pi$	$(\sigma_{\theta_P^{(1)}}, \sigma_{\theta_A^{(1)}}) / \pi$
ISRES	(0.89, 0.98)	(0.07, 0.53)
STOGO	(0.73, 1.0)	(≈ 0 , ≈ 0)
MLSL + BOBQYA	(0.81, 1.88)	(0.04, 0.27)
MLSL + LBFGS	(0.5, 1.0)	(0.35, 0.71)
MLSL + MMA	(0.81, 0.14)	(0.21, 0.04)
SAMIN	(0.87, 0.99)	(0.07, 0.59)
NTR	(0.5, 1.0)	(0.5, 1.0)
APS	(1.0, 1.04)	(≈ 0 , 0.54)
IPOPT	(0.45, 0.89)	(0.26, 0.57)
Optimal Values	(1, -)	(0, -)

(a) Inferred angles for neutrino 1

Method	$(\mu_{\theta_P^{(2)}}, \mu_{\theta_A^{(2)}}) / \pi$	$(\sigma_{\theta_P^{(2)}}, \sigma_{\theta_A^{(2)}}) / \pi$
ISRES	(0.17, 0.98)	(0.10, 0.53)
STOGO	(0.43, 0.99)	(≈ 0 , ≈ 0)
MLSL + BOBQYA	(0.15, 0.18)	(0.03, 0.26)
MLSL + LBFGS	(0.5, 1.0)	(0.35, 0.71)
MLSL + MMA	(0.21, 0.33)	(0.04, 0.22)
SAMIN	(0.19, 1.0)	(0.07, 0.59)
NTR	(0.35, 0.71)	(0.35, 0.71)
APS	(≈ 0 , 1.03)	(≈ 0 , 0.64)
IPOPT	(0.19, 0.89)	(0.21, 0.59)
Optimal Values	(0, -)	(0, -)

(b) Inferred angles for neutrino 2

TABLE III. Sample average and standard deviation of the inferred angles in flavor space, for the initial polarization vectors of the **two** neutrinos at the end of the each optimization procedure. The optimal values shown at the end of the table; there is no preferred azimuthal angle. These results are for the set of experiments described in the caption of Fig. 11.

scenario, for the detector data generated based on table I, we remove the neutrino self-interaction from Eq. 3 and optimize for the matter coupling potential V_0 . In Fig. 12 we simultaneously plot the final inferred value of V_0 by the APS method, and its cost value, as function of the initial guess. To obtain this plots we used a grid of 500 uniform initial guesses in the range $[0, 1000]$ and 1000 optimization iterations were employed per initial value. As the plot shows, there is no convergence to a global minimum, and the smallest cost value obtained is about five orders of magnitude higher than the one found by the same optimization method is Sec. III. These are indications of how poorly the current model performs in reproducing the detector data. One could expect such differences to be used in discerning among the various proposed models when a detection is made.

B. Inclusion of experimental uncertainties

A second point that merits consideration is the inclusion of experimental uncertainties in inference. The results presented here can be interpreted as maximum likelihood estimates obtained from using the central value of measured neutrino polarization vectors. Assuming the

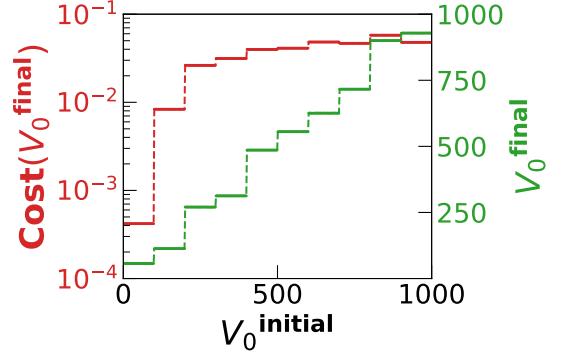


FIG. 12. The final inferred value of V_0 (green) with the APS optimization method, and its corresponding cost value (red), as function of the initial guess V_0^{initial} . In this experiment, the optimization procedure was given a model with the neutrino-neutrino coupling $\mu_0 = 0$, even though the simulated “detector” data was generated using a model with $\mu_0 = 10$. The lack of a global convergence and the relatively high values of the cost function are indications that the model with $\mu_0 = 0$ does not properly reproduce the detector measurements.

measurement error to be tiny one can approximate,

$$\frac{\Delta \text{Cost}}{\Delta V_0} \approx \frac{\partial \text{Cost}}{\partial V_0}, \quad (6)$$

and estimate the error in the inferred value of V_0 . In the rare event the optimal solution is found, $\frac{\partial \text{Cost}}{\partial V_0} = 0$, and the local gradient provides no useful information in estimating uncertainties. In practice, the optimization procedure converges close to the optimal solutions so local gradients are non-vanishing. However, detection errors may be small but not tiny and the outlined procedure might not provide robust estimates. Determining parameter uncertainties is, indeed, a rather complicated task and the presence of multiple minima with a cost value very close to the global optimum could lead to situations in which a small uncertainty in measurement could lead to large uncertainties in the parameter estimation. This is why one would need to probe large regions of parameter space through Monte Carlo sampling to properly represent uncertainties. When there is some information on the parameters of the model under consideration, one could perform a Bayesian analysis. These directions deserve to be treated in depth and are outside the scope of this work. We plan to pursue them in the future.

VIII. CONCLUSION

We have combined recent developments in deep learning with data assimilation, to examine what information is contained within a detected neutrino signal regarding complex astrophysical environments such as supernovae.

By recasting the differential equations of our model as layers of a neural architecture, we accomplish two things.

Firstly, we automatically satisfy physical constraints for the problem under investigation. Secondly, we demonstrate a model framework in which domain discretization is absent. This framework has allowed us to focus on the prediction error (that is, the cost function), by eliminating any free parameters associated with the grid resulting from domain discretization.

In addition, we have tested nine optimization algorithms that cover a wide range of techniques, and we have identified the “Adaptive Particle Swarm Algorithm” (APS) as best suited for our purposes in this paper. This algorithm, through its four evolutionary stages, is able to move out of local minima and thus has a high chance of finding a global minimum.

The study conducted here has focused on a small system, primarily as a first testing ground for our framework. We expect the computational complexity to inevitably increase with larger systems, and many degeneracies to be present in the parameter space. Thus, when the particle number is greatly increased, we might combine evolutionary algorithms such as APS for a wide parameter search, with a follow-up gradient- and Hessian-based method such as IPOPT as a secondary search within smaller optimal regions that are found by the first search. We might also need to transition to distributed ordinary differential equation solvers, which can take advantage of computer clusters.

We intend to maintain a level of modeling complexity lower than that of three dimensional supernovae simulations (which take months for a single run to complete), and provide a computational service that is complementary to simulations and can function as a bridge to earth-based neutrino detection. A more realistic setting, however, would require more than one affine parameter: temporal, in addition to spatial evolution. We may need to make away with rotational symmetry, leading to three affine spatial parameters. In this case, a combined approach, with the grid discretization of previous work applied to the spatial dimensions, and the new framework developed here applied to the temporal parameter, might prove useful.

ACKNOWLEDGMENTS

We thank Anthony Mezzacappa for the useful discussions. GMF acknowledges NSF Grant No. PHY-1914242 at UCSD. ER, AVP, and GMF acknowledge the NSF N3AS Physics Frontier Center, NSF Grant No. PHY-2020275, and the Heising-Simons Foundation (2017-228). EA acknowledges an Institutional Support for Research and Creativity grant from New York Institute of Technology. Additionally, GMF acknowledges Department of Energy Scientific Discovery through Advanced Computing (SciDAC-4) grant register No. SN60152 (award number de-sc0018297). The work of AVP was also supported in part by the U.S. Department of Energy under contract number DE-AC02-76SF00515. The authors acknowledge

the Minnesota Supercomputing Institute (MSI) at the University of Minnesota for providing resources that contributed to the research results reported within this paper. See <http://www.msi.umn.edu>. In addition, HSL, a collection of Fortran codes for large scale scientific computation was used in conjunction with IPOPT. See <http://www.hsl.rl.ac.uk>.

-
- [1] A Burrows, “Neutrinos from supernova explosions,” *Annual Review of Nuclear and Particle Science* **40**, 181–212 (1990).
- [2] Alessandro Mirizzi, Irene Tamborra, Hans-Thomas Janka, Ninetta Saviano, Kate Scholberg, Robert Bollig, Lorenz Hudepohl, and Sovan Chakraborty, “Supernova Neutrinos: Production, Oscillations and Detection,” *Riv. Nuovo Cim.* **39**, 1–112 (2016).
- [3] Gabriel Martínez-Pinedo, Tobias Fischer, Karlheinz Langanke, Andreas Lohs, Andre Sieverding, and Meng-Ru Wu, “Neutrinos and their impact on core-collapse supernova nucleosynthesis,” in *Handbook of Supernovae*, edited by Athem W. Alsabti and Paul Murdin (Springer International Publishing, Cham, 2017) pp. 1805–1841.
- [4] B. Müller, “Neutrino emission as diagnostics of core-collapse supernovae,” *Annual Review of Nuclear and Particle Science* **69**, 253–278 (2019).
- [5] Huaiyu Duan, George M. Fuller, J Carlson, and Yong-Zhong Qian, “Simulation of Coherent Non-Linear Neutrino Flavor Transformation in the Supernova Environment. 1. Correlated Neutrino Trajectories,” *Phys. Rev. D* **74**, 105014 (2006), arXiv:astro-ph/0606616.
- [6] Huaiyu Duan, George M. Fuller, and Yong-Zhong Qian, “Collective neutrino flavor transformation in supernovae,” *Phys. Rev. D* **74**, 123004 (2006), arXiv:astro-ph/0511275.
- [7] Huaiyu Duan, George M. Fuller, J. Carlson, and Yong-Zhong Qian, “Coherent Development of Neutrino Flavor in the Supernova Environment,” *Phys. Rev. Lett.* **97**, 241101 (2006), arXiv:astro-ph/0608050.
- [8] Steen Hannestad, Georg G. Raffelt, Gunter Sigl, and Yvonne Y.Y. Wong, “Self-induced conversion in dense neutrino gases: Pendulum in flavour space,” *Phys. Rev. D* **74**, 105010 (2006), [Erratum: Phys.Rev.D 76, 029901 (2007)], arXiv:astro-ph/0608695.
- [9] Georg G. Raffelt and Alexei Yu. Smirnov, “Self-induced spectral splits in supernova neutrino fluxes,” *Phys. Rev. D* **76**, 081301 (2007), [Erratum: Phys.Rev.D 77, 029903 (2008)], arXiv:0705.1830 [hep-ph].
- [10] Georg G. Raffelt and Alexei Yu. Smirnov, “Adiabaticity and spectral splits in collective neutrino transformations,” *Phys. Rev. D* **76**, 125008 (2007), arXiv:0709.4641 [hep-ph].
- [11] Basudeb Dasgupta and Amol Dighe, “Collective three-flavor oscillations of supernova neutrinos,” *Phys. Rev. D* **77**, 113002 (2008), arXiv:0712.3798 [hep-ph].
- [12] Lucas Johns and George M. Fuller, “Strange mechanics of the neutrino flavor pendulum,” *Phys. Rev. D* **97**, 023020 (2018), arXiv:1709.00518 [hep-ph].
- [13] John F. Cherry, J. Carlson, Alexander Friedland, George M. Fuller, and Alexey Vlasenko, “Halo Modification of a Supernova Neutronization Neutrino Burst,” *Phys. Rev. D* **87**, 085037 (2013), arXiv:1302.1159 [astro-ph.HE].
- [14] Masamichi Zaizen, John F. Cherry, Tomoya Takiwaki, Shunsaku Horiuchi, Kei Kotake, Hideyuki Umeda, and Takashi Yoshida, “Neutrino halo effect on collective neutrino oscillation in iron core-collapse supernova model of a $9.6 M_{\odot}$ star,” *JCAP* **06**, 011 (2020), arXiv:1908.10594 [astro-ph.HE].
- [15] Huaiyu Duan and James P Kneller, “Neutrino flavour transformation in supernovae,” *J. Phys. G* **36**, 113201 (2009), arXiv:0904.0974 [astro-ph.HE].
- [16] Huaiyu Duan, George M. Fuller, and Yong-Zhong Qian, “Collective Neutrino Oscillations,” *Ann. Rev. Nucl. Part. Sci.* **60**, 569–594 (2010), arXiv:1001.2799 [hep-ph].
- [17] Alessandro Mirizzi, Irene Tamborra, Hans-Thomas Janka, Ninetta Saviano, Kate Scholberg, Robert Bollig, Lorenz Hudepohl, and Sovan Chakraborty, “Supernova Neutrinos: Production, Oscillations and Detection,” *Riv. Nuovo Cim.* **39**, 1–112 (2016), arXiv:1508.00785 [astro-ph.HE].
- [18] Sovan Chakraborty, Rasmus Hansen, Ignacio Izaguirre, and Georg Raffelt, “Collective neutrino flavor conversion: Recent developments,” *Nucl. Phys. B* **908**, 366–381 (2016), arXiv:1602.02766 [hep-ph].
- [19] Arka Banerjee, Amol Dighe, and Georg Raffelt, “Linearized flavor-stability analysis of dense neutrino streams,” *Phys. Rev. D* **84**, 053013 (2011), arXiv:1107.2308 [hep-ph].
- [20] Georg Raffelt, Srdjan Sarikas, and David de Sousa Seixas, “Axial Symmetry Breaking in Self-Induced Flavor Conversion of Supernova Neutrino Fluxes,” *Phys. Rev. Lett.* **111**, 091101 (2013), [Erratum: Phys.Rev.Lett. 113, 239903 (2014)], arXiv:1305.7140 [hep-ph].
- [21] Sovan Chakraborty and Alessandro Mirizzi, “Multi-azimuthal-angle instability for different supernova neutrino fluxes,” *Phys. Rev. D* **90**, 033004 (2014), arXiv:1308.5255 [hep-ph].
- [22] Huaiyu Duan and Shashank Shalgar, “Flavor instabilities in the neutrino line model,” *Phys. Lett. B* **747**, 139–143 (2015), arXiv:1412.7097 [hep-ph].
- [23] Sajad Abbar and Huaiyu Duan, “Neutrino flavor instabilities in a time-dependent supernova model,” *Physics Letters B* **751**, 43–47 (2015), arXiv:1509.01538 [astro-ph.HE].
- [24] S. Chakraborty, R. S. Hansen, I. Izaguirre, and G. G. Raffelt, “Self-induced flavor conversion of supernova neutrinos on small scales,” *JCAP* **2016**, 028 (2016), arXiv:1507.07569 [hep-ph].
- [25] R. F. Sawyer, “Multiangle instability in dense neutrino systems,” *Phys. Rev. D* **79**, 105003 (2009), arXiv:0803.4319 [astro-ph].
- [26] R. F. Sawyer, “Neutrino Cloud Instabilities Just above the Neutrino Sphere of a Supernova,” *Phys. Rev. Lett.* **116**, 081101 (2016), arXiv:1509.03323 [astro-ph.HE].
- [27] Basudeb Dasgupta and Alessandro Mirizzi, “Temporal instability enables neutrino flavor conversions deep inside supernovae,” *Phys. Rev. D* **92**, 125030 (2015), arXiv:1509.03171 [hep-ph].
- [28] Ignacio Izaguirre, Georg Raffelt, and Irene Tamborra, “Fast Pairwise Conversion of Supernova Neutrinos: A Dispersion-Relation Approach,” *Phys. Rev. Lett.* **118**, 021101 (2017), arXiv:1610.01612 [hep-ph].
- [29] Basudeb Dasgupta, Alessandro Mirizzi, and Manibrata Sen, “Fast neutrino flavor conversions near the supernova core with realistic flavor-dependent angular distributions,” *JCAP* **02**, 019 (2017), arXiv:1609.00528 [hep-ph].
- [30] Francesco Capozzi, Basudeb Dasgupta, Eligio Lisi, Antonio Marrone, and Alessandro Mirizzi, “Fast flavor

- conversions of supernova neutrinos: Classifying instabilities via dispersion relations,” *Phys. Rev. D* **96**, 043016 (2017), [arXiv:1706.03360 \[hep-ph\]](#).
- [31] Basudeb Dasgupta and Manibrata Sen, “Fast neutrino flavor conversion as oscillations in a quartic potential,” *Phys. Rev. D* **97**, 023017 (2018), [arXiv:1709.08671 \[hep-ph\]](#).
- [32] Sajad Abbar and Maria Cristina Volpe, “On fast neutrino flavor conversion modes in the nonlinear regime,” *Physics Letters B* **790**, 545–550 (2019), [arXiv:1811.04215 \[astro-ph.HE\]](#).
- [33] Shashank Shalgar and Irene Tamborra, “On the Occurrence of Crossings between the Angular Distributions of Electron Neutrinos and Antineutrinos in the Supernova Core,” *Astrophys. J.* **883**, 80 (2019), [arXiv:1904.07236 \[astro-ph.HE\]](#).
- [34] Francesco Capozzi, Basudeb Dasgupta, Alessandro Mirizzi, Manibrata Sen, and Günter Sigl, “Collisional Triggering of Fast Flavor Conversions of Supernova Neutrinos,” *Phys. Rev. Lett.* **122**, 091101 (2019), [arXiv:1808.06618 \[hep-ph\]](#).
- [35] Milad Delfan Azari, Shoichi Yamada, Taiki Morinaga, Wakana Iwakami, Hirotada Okawa, Hiroki Nagakura, and Kohsuke Sumiyoshi, “Linear analysis of fast-pairwise collective neutrino oscillations in core-collapse supernovae based on the results of Boltzmann simulations,” *Phys. Rev. D* **99**, 103011 (2019), [arXiv:1902.07467 \[astro-ph.HE\]](#).
- [36] Hiroki Nagakura, Taiki Morinaga, Chinami Kato, and Shoichi Yamada, “Fast-pairwise Collective Neutrino Oscillations Associated with Asymmetric Neutrino Emissions in Core-collapse Supernovae,” *Astrophys. J.* **886**, 139 (2019), [arXiv:1910.04288 \[astro-ph.HE\]](#).
- [37] Lucas Johns, Hiroki Nagakura, George M. Fuller, and Adam Burrows, “Neutrino oscillations in supernovae: angular moments and fast instabilities,” *Phys. Rev. D* **101**, 043009 (2020), [arXiv:1910.05682 \[hep-ph\]](#).
- [38] Madhurima Chakraborty and Sovan Chakraborty, “Three flavor neutrino conversions in supernovae: slow & fast instabilities,” *JCAP* **01**, 005 (2020), [arXiv:1909.10420 \[hep-ph\]](#).
- [39] Taiki Morinaga and Shoichi Yamada, “Linear stability analysis of collective neutrino oscillations without spurious modes,” *Phys. Rev. D* **97**, 023024 (2018), [arXiv:1803.05913 \[hep-ph\]](#).
- [40] Sajad Abbar, Huaiyu Duan, Kohsuke Sumiyoshi, Tomoya Takiwaki, and Maria Cristina Volpe, “On the occurrence of fast neutrino flavor conversions in multidimensional supernova models,” *Phys. Rev. D* **100**, 043004 (2019), [arXiv:1812.06883 \[astro-ph.HE\]](#).
- [41] John F. Cherry, George M. Fuller, Shunsaku Horiuchi, Kei Kotake, Tomoya Takiwaki, and Tobias Fischer, “Time of Flight and Supernova Progenitor Effects on the Neutrino Halo,” (2019), [arXiv:1912.11489 \[astro-ph.HE\]](#).
- [42] Taiki Morinaga, Hiroki Nagakura, Chinami Kato, and Shoichi Yamada, “Fast neutrino-flavor conversion in the preshock region of core-collapse supernovae,” *Phys. Rev. Res.* **2**, 012046 (2020), [arXiv:1909.13131 \[astro-ph.HE\]](#).
- [43] Soumya Bhattacharyya and Basudeb Dasgupta, “Late-time behavior of fast neutrino oscillations,” *Phys. Rev. D* **102**, 063018 (2020), [arXiv:2005.00459 \[hep-ph\]](#).
- [44] Sajad Abbar, “Searching for Fast Neutrino Flavor Conversion Modes in Core-collapse Supernova Simulations,” *JCAP* **05**, 027 (2020), [arXiv:2003.00969 \[astro-ph.HE\]](#).
- [45] Francesco Capozzi, Sajad Abbar, Robert Bollig, and H. Thomas Janka, “Fast neutrino flavor conversions in one-dimensional core-collapse supernova models with and without muon creation,” (2020), [arXiv:2012.08525 \[astro-ph.HE\]](#).
- [46] Irene Tamborra and Shashank Shalgar, “New Developments in Flavor Evolution of a Dense Neutrino Gas,” (2020), [10.1146/annurev-nucl-102920-050505, arXiv:2011.01948 \[astro-ph.HE\]](#).
- [47] Babak Abi *et al.* (DUNE), “Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II DUNE Physics,” (2020), [arXiv:2002.03005 \[hep-ex\]](#).
- [48] K. Abe *et al.* (Hyper-Kamiokande), “Hyper-Kamiokande Design Report,” (2018), [arXiv:1805.04163 \[physics.ins-det\]](#).
- [49] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016).
- [50] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le, “Learning transferable architectures for scalable image recognition,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018).
- [51] Boukaye Boubacar Traore, Bernard Kamsu-Foguem, and Fana Tangara, “Deep convolution neural network for image recognition,” *Ecol. Informatics* **48**, 257–268 (2018).
- [52] Tom Young, Devamanyu Hazarika, Soujanya Poria, and Erik Cambria, “Recent trends in deep learning based natural language processing,” *CoRR* **abs/1708.02709** (2017), [1708.02709](#).
- [53] Daniel W. Otter, Julian R. Medina, and Jugal K. Kalita, “A survey of the usages of deep learning in natural language processing,” *CoRR* **abs/1807.10854** (2018).
- [54] Binhua Tang, Zixiang Pan, Kang Yin, and Asif Khateeb, “Recent advances of deep learning in bioinformatics and computational biology,” *Frontiers in genetics* **10**, 214–214 (2019).
- [55] Ryad Zemouri, Nouredine Zerhouni, and Daniel Racocanu, “Deep learning in the biomedical applications: Recent and future status,” *Applied Sciences* **9**, 1526 (2019).
- [56] Eve Armstrong, Amol V. Patwardhan, Lucas Johns, Chad T. Kishimoto, Henry D. I. Abarbanel, and George M. Fuller, “An optimization-based approach to calculating neutrino flavor evolution,” *Phys. Rev. D* **96**, 083008 (2017).
- [57] Eve Armstrong, Amol V. Patwardhan, Ermal Rrapaj, Sina Fallah Ardizi, and George M. Fuller, “Inference offers a metric to constrain dynamical models of neutrino flavor transformation,” *Phys. Rev. D* **102**, 043013 (2020).
- [58] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud, “Neural ordinary differential equations,” in *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., 2018) pp. 6571–6583.
- [59] Kurt Hornik, Maxwell Stinchcombe, and Halbert White, “Multilayer feedforward networks are universal approximators,” *Neural Networks* **2**, 359 – 366 (1989).
- [60] B.D. Keister, “Numerical and Physical Stability of Supernova Neutrino Flavor Evolution,” *Phys. Scripta* **90**, 088008 (2015), [arXiv:1408.4729 \[astro-ph.HE\]](#).

- [61] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation*, Other Titles in Applied Mathematics (Society for Industrial and Applied Mathematics, 2005).
- [62] BP vanMilligen, Tribaldos, and Jiménez, “Neural network differential equation and plasma equilibrium solver.” *Physical review letters* **75** **20**, 3594–3597 (1995).
- [63] I. E. Lagaris, A. Likas, and D. I. Fotiadis, “Artificial neural networks for solving ordinary and partial differential equations,” *IEEE Transactions on Neural Networks* **9**, 987–1000 (1998).
- [64] Jens Berg and Kaj Nyström, “A unified deep artificial neural network approach to partial differential equations in complex geometries,” *Neurocomputing* **317**, 28–41 (2018).
- [65] Martin Magill, Faisal Qureshi, and Hendrick de Haan, “Neural networks trained to solve differential equations learn general representations,” in *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., 2018) pp. 4071–4081.
- [66] M. Raissi, P. Perdikaris, and G.E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics* **378**, 686 – 707 (2019).
- [67] Ernst Hairer and Gerhard Wanner, “Stiff differential equations solved by radau methods,” *Journal of Computational and Applied Mathematics* **111**, 93 – 111 (1999).
- [68] J. Revels, M. Lubin, and T. Papamarkou, “Forward-mode automatic differentiation in Julia,” [arXiv:1607.07892 \[cs.MS\]](https://arxiv.org/abs/1607.07892) (2016).
- [69] Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind, “Automatic differentiation in machine learning: a survey,” *Journal of Machine Learning Research* **18**, 1–43 (2018).
- [70] Adrian Sandu, Dacian N. Daescu, and Gregory R. Carmichael, “Direct and adjoint sensitivity analysis of chemical kinetic systems with kpp: Part i—theory and software tools,” *Atmospheric Environment* **37**, 5083 – 5096 (2003).
- [71] Yang Cao, Shengtai Li, Linda Petzold, and Radu Serban, “Adjoint sensitivity analysis for differential-algebraic equations: The adjoint dae system and its numerical solution,” *SIAM Journal on Scientific Computing* **24**, 1076–1089 (2003).
- [72] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah, “Julia: A fresh approach to numerical computing,” *SIAM Review* **59**, 65–98 (2017).
- [73] Christopher Rackauckas and Q Nie, “Differential equations.jl – a performant and feature-rich ecosystem for solving differential equations in julia,” *Journal of Open Research Software* **5**, 15 (2017).
- [74] Michael Innes, Elliot Saba, Keno Fischer, Dhairya Gandhi, Marco Concetto Rudilosso, Neethu Mariya Joy, Tejan Karmali, Avik Pal, and Viral Shah, “Fashionable modelling with flux,” *CoRR* **abs/1811.01457** (2018).
- [75] Christopher Rackauckas, Mike Innes, Yingbo Ma, Jesse Bettencourt, Lyndon White, and Vaibhav Dixit, “DiffEqFlux.jl - A julia library for neural differential equations,” *CoRR* **abs/1902.02376** (2019).
- [76] Léon Bottou, “Online algorithms and stochastic approximations,” in *Online Learning and Neural Networks*, edited by David Saad (Cambridge University Press, Cambridge, UK, 1998).
- [77] John Duchi, Elad Hazan, and Yoram Singer, “Adaptive subgradient methods for online learning and stochastic optimization,” *J. Mach. Learn. Res.* **12**, 2121–2159 (2011).
- [78] Diederik P Kingma and Jimmy Lei Ba, “Adam: A method for stochastic gradient descent,” in *ICLR: International Conference on Learning Representations* (2015).
- [79] Robert Malouf, “A comparison of algorithms for maximum entropy parameter estimation,” in *Proceedings of the 6th Conference on Natural Language Learning - Volume 20*, COLING-02 (Association for Computational Linguistics, USA, 2002) p. 1–7.
- [80] William Goffe, Gary Ferrier, and John Rogers, “Global optimization of statistical functions with simulated annealing,” *Journal of Econometrics* **60**, 65–99 (1994).
- [81] Goffe William L., “Simann: A global optimization algorithm using simulated annealing,” *Studies in Nonlinear Dynamics & Econometrics* **1**, 1–9 (1996).
- [82] T. P. Runarsson and Xin Yao, “Stochastic ranking for constrained evolutionary optimization,” *IEEE Transactions on Evolutionary Computation* **4**, 284–294 (2000).
- [83] T. P. Runarsson and Xin Yao, “Search biases in constrained evolutionary optimization,” *Trans. Sys. Man Cyber Part C* **35**, 233–243 (2005).
- [84] John A. Nelder and Roger Mead, “A simplex method for function minimization,” *Computer Journal* **7**, 308–313 (1965).
- [85] Z. Zhan, J. Zhang, Y. Li, and H. S. Chung, “Adaptive particle swarm optimization,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* **39**, 1362–1381 (2009).
- [86] Andreas Wächter and Lorenz T. Biegler, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Math. Program.* **106**, 25–57 (2006).
- [87] Jorge Nocedal and Stephen J. Wright, *Numerical Optimization*, 2nd ed. (Springer, New York, NY, USA, 2006).
- [88] Kaj Madsen and Julius Žilinskas, “Parallel branch-and-bound attraction based methods for global optimization,” in *Stochastic and Global Optimization*, edited by Gintautas Dzemyda, Vydūnas Šaltenis, and Antanas Žilinskas (Springer US, Boston, MA, 2002) pp. 175–187.
- [89] A. H. G. Rinnooy Kan and G. T. Timmer, “Stochastic global optimization methods part i: Clustering methods,” *Mathematical Programming* **39**, 27–56 (1987).
- [90] Sergei Kucherenko and Yury Sytsko, “Application of deterministic low-discrepancy sequences in global optimization,” *Computational Optimization and Applications* **30**, 297–318 (2005).
- [91] M. J. D. Powell, “The bobyqa algorithm for bound constrained optimization without derivatives,” technical report NA2009/06 (2009).
- [92] Krister Svanberg, “A class of globally convergent optimization methods based on conservative convex separable approximations,” *SIAM Journal on Optimization* **12**, 555–573 (2002).
- [93] Jorge Nocedal, “Updating quasi-newton matrices with limited storage,” *Mathematics of Computation* **35**, 773–782 (1980).
- [94] Dong C. Liu and Jorge Nocedal, “On the limited memory bfgs method for large scale optimization,” *Mathematical Programming* **45**, 503–528 (1989).

- [95] David G. Luenberger and Yinyu Ye, *Linear and Nonlinear Programming* (Springer Publishing Company, Incorporated, 2015).
- [96] P. A. Vikhar, “Evolutionary algorithms: A critical review and its future prospects,” in *2016 International Conference on Global Trends in Signal Processing, Information Computing and Communication (ICGTSPICC)* (2016) pp. 261–265.
- [97] Steven G. Johnson, “The nlopt nonlinear optimization package,” ab-initio.mit.edu/nlopt.
- [98] Patrick Kofod Mogensen and Asbjørn Nilsen Riseth, “Optim: A mathematical optimization package for Julia,” *Journal of Open Source Software* **3**, 615 (2018).