Physics-Informed Neural Networks for Derivative-Constrained PDEs

Kentaro Hoshisashi [1] **Carolyn E. Phelan** [1] **Paolo Barucca** [1]

Abstract

Physics-Informed Neural Networks (PINNs) have emerged as a promising approach for solving partial differential equations (PDEs) using deep learning. However, standard PINNs do not address the problem of constrained PDEs, where the solution must satisfy additional equality or inequality constraints beyond the governing equations. In this paper, we introduce Derivative-Constrained PINNs (DC-PINNs), a novel framework that seamlessly incorporates constraint information into the PINNs training process. DC-PINNs employ a constraint-aware loss function that penalizes constraint violations while simultaneously minimizing the PDE residual. Key components include self-adaptive loss balancing techniques that automatically tune the relative weighting of each term, enhancing training stability, and the use of automatic differentiation to efficiently compute exact derivatives. This study demonstrates the effectiveness of DC-PINNs on several benchmark problems related to quantitative finance: heat diffusion, Black-Scholes pricing, and local volatility surface calibration. The results showcase improvements in generating appropriate solutions that satisfy the constraints compared to baseline PINNs methods. The DC-PINNs framework opens up new possibilities for solving constrained PDEs in multi-objective optimization.

1. Introduction

Partial differential equations (PDEs) play a crucial role in modelling various physical phenomena across scientific and engineering disciplines. Traditional numerical methods for solving PDEs, such as finite difference and finite element methods, have been widely used but often face challenges in terms of computational efficiency and handling complex geometries. Recently, Physics-Informed Neural Networks (PINNs) by (Raissi et al., 2019) have emerged as a promising alternative, leveraging deep learning to solve PDEs with high accuracy and efficiency.

However, despite their empirical successes, PINNs still face challenges in the presence of additional equality or inequality constraints beyond the PDE itself, as pioneered in (Lagaris et al., 1998). Such constrained PDEs are ubiquitous in real-world applications: heat diffusion with temperature bounds, option pricing with no-arbitrage constraints, and fluid flow with velocity limits, to name a few. Naive techniques for embedding constraints can lead to unstable training, slow convergence, and constraint violation. Although more sophisticated approaches have been proposed to better handle constraints in PINNs, limitations persist. Conservative PINNs (Jagtap et al., 2020) and DC NN (Lo & Huang, 2023) enforce equality constraints, while Augmented Lagrangian approaches (Lu et al., 2021) and theory-guided neural networks (Chen et al., 2021) show stronger performance on inequality-constrained PDEs. Nevertheless, there is still a need for methods that can effectively handle constraints while reducing the dependence on intricate hyperparameter adjustments and problem-specific architectures.

In this study, we propose Derivative-Constrained PINNs (DC-PINNs), a general and robust framework for solving derivative-constrained PDEs using deep learning. Our approach seamlessly integrates the constraints into the learning process, ensuring that the solution satisfies the prescribed conditions while maintaining the benefits of PINNs. The key contributions of this study include:

- A flexible constraint-aware loss function that admits general nonlinear constraints and seamlessly incorporates them into the PDE residual objective.
- Self-adaptive loss balancing techniques that automatically tune the weightings of the objective terms, including derivatives obtained through automatic differentiation stabilizing training across diverse problem settings.
- Demonstration of the approximation ability of DC-PINNs on benchmark PDEs from the heat equations, structural modelling of finance, and its inverse problem, showcasing improvements over PINNs approaches.

[*]Equal contribution [1]Department of Computer Science, University College London, London, United Kingdom. Correspondence to: Kentaro Hoshisashi <k.hoshisashi@ucl.ac.uk>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

2. Problem Formulation

2.1. Derivative-Constrained PDE Problem

Physics-based machine learning can be formulated as an optimization problem that aims to find the feasible parameters $\hat{\theta}$ of a parameterized function $y := \varphi_\theta(x)$ while satisfying constraints represented by PDEs and related conditions by

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}}\ \mathcal{L}(x, \mathcal{D}y), \tag{1}$$

$$\begin{aligned} \text{s.t.} \qquad & f(x, \mathcal{D}y) = 0, \quad \text{governing PDEs} \\ & \mathcal{B} = \left\{ b_k, \mathcal{D}_k^{(b)} | \mathcal{D}_k^{(b)} y(x) = 0, x \in b_k \right\}_{k=1}^{N_b} \\ & \mathcal{H} = \left\{ h_k, \mathcal{D}_k^{(h)} | \mathcal{D}_k^{(h)} y(x) \geq 0, x \in h_k \right\}_{k=1}^{N_h} \end{aligned} \tag{2}$$

where $\mathcal{L}$ is the objective loss function, multivariate inputs $x \in \mathbb{R}^n$ are state variables, $\mathcal{D}$ is the set of individual differential operators including the 0th order, f represents the governing PDEs, $\mathcal{B}$ is the set of boundary equality constraints, $\mathcal{H}$ is the set of inequality constraints, and $N_{(\cdot)}$ denotes the number of corresponding conditions.

This study focuses on inequality constraints involving derivatives and employs a discretize-then-optimize approach combined with gradient-based optimization using artificial neural networks (ANNs). The optimization problem is first discretized numerically, transforming it into a finite-dimensional problem. In the ANN approach, automatic differentiation can be used to compute derivatives of the network output with respect to the input variables. However, the standard ANN architecture does not inherently satisfy all PDEs and derivative conditions. Therefore, additional techniques or modifications may be necessary to ensure that the ANN solution complies with the governing equations and constraints.

2.2. Physics-Informed Neural Networks (PINNs)

Physics-Informed Neural Networks (PINNs), introduced in (Raissi et al., 2019), are neural networks that incorporate underlying physical laws into the architecture through PDEs, forming a new class of data-efficient universal function approximations. We consider a parametrized PDE system given by

$$\begin{aligned} \boldsymbol{f}[\varphi(x)], x \in \Omega, \\ \boldsymbol{b}[\varphi(x)], x \in \partial\Omega, \end{aligned} \tag{3}$$

where $\boldsymbol{f}$ and $\boldsymbol{b}$ are a set of PDE and boundary operators, Ω and $\partial\Omega$ are the spatial domain and the boundary.

PINNs solve this PDE system as an optimization problem using an artificial neural network by minimizing the total loss in a deep learning context:

$$\mathcal{L} := \mathcal{L}_0 + \mathcal{L}_b + \mathcal{L}_f, \tag{4}$$

$$\begin{aligned} \mathcal{L}_0 &:= \frac{1}{N_0} \sum_{i=1}^{N_0} \left| \varphi(x_0^{(i)}) - \hat{y}_0^{(i)} \right|^2, \\ \mathcal{L}_b &:= \frac{1}{N_b} \sum_{i=1}^{N_b} \left| b[\varphi(x_b^{(i)})] \right|^2, \\ \mathcal{L}_f &:= \frac{1}{N_f} \sum_{i=1}^{N_f} \left| f[\varphi(x_f^{(i)})] \right|^2, \end{aligned} \tag{5}$$

where $\mathcal{L}_0$ represents the error between observed ($\hat{y}$) and predicted values, $\mathcal{L}_b$ enforces boundary conditions, and $\mathcal{L}_f$ penalizes the PDE residual at a set of collocation points.

2.3. Derivative-Constrained PINNs (DC-PINNs)

We now consider the extension of PINNs to handle derivative inequality constraints, which we term Derivative-Constrained PINNs (DC-PINNs). Assume the presence of inequality constraints of the form

$$\boldsymbol{h}[\varphi(x)], x \in \Omega \tag{6}$$

where $\boldsymbol{h}[\cdot]$ represents a set of differential operators acting on an inequality equation, which includes derivatives. There are several methods available to enforce inequality conditions in general. The direct approach is to formulate loss functions of inequalities and impose them as soft constraints with fixed loss weights. To fit with inequality constraints, a loss $\mathcal{L}_h$ to be minimized is defined as

$$\mathcal{L}_h := \frac{1}{N_h} \sum_{i=1}^{N_h} \gamma \circ \left| h\left(\varphi(x_h^{(i)})\right) \right|^2. \tag{7}$$

$$\gamma(x) = \begin{cases} x, & \text{if inequality is not satisfied} \\ 0, & \text{otherwise} \end{cases} \tag{8}$$

where $\mathcal{L}_h$ penalizes the violation of inequality constraints at a set of collocation points, and the function γ determines the penalty based on whether the inequality is satisfied or not.

Our proposed framework extends PINNs to handle derivative-constrained PDEs effectively by seamlessly integrating the constraints into the learning process. We introduce a constraint-aware loss function that penalizes violations of the constraints while simultaneously minimizing the residual loss. However, setting large loss weights can cause an ill-conditioned problem. On the other hand, when small loss weights are chosen, the estimated solution may violate the inequalities. In this sense, we formulate the total cost in DC-PINNs to be minimized as,

$$\mathcal{L} := \lambda_0 \hat{\mathcal{L}}_0 + \lambda_b \hat{\mathcal{L}}_b + \lambda_f \hat{\mathcal{L}}_f + \lambda_h \hat{\mathcal{L}}_h, \tag{9}$$

where λ are weighting coefficients for each categorized loss term. By minimizing this total loss, the neural network

approximation satisfies the governing PDE, boundary/initial conditions, and the prescribed inequality constraints. The categorized loss terms are defined as

$$\hat{\mathcal{L}}_0 = \frac{1}{N_0}\sum_{i=1}^{N_0} m_0^{(i)} \left|\varphi_\theta\left(x_0^{(i)}\right) - y_0^{(i)}\right|^2, \tag{10}$$

$$\hat{\mathcal{L}}_b = \frac{1}{N_b}\sum_{i=1}^{N_b} m_b^{(i)} \left|\varphi_\theta\left(x_b^{(i)}\right) - y_b^{(i)}\right|^2, \tag{11}$$

$$\hat{\mathcal{L}}_f = \frac{1}{N_f}\sum_{i=1}^{N_f} m_f^{(i)} \left|f\left(\varphi_\theta(x_f^{(i)})\right)\right|^2, \tag{12}$$

$$\hat{\mathcal{L}}_h := \frac{1}{N_h}\sum_{i=1}^{N_h} m_h^{(i)} \left|\gamma \circ h\left(\varphi_\theta(x_h^{(i)})\right)\right|^2, \tag{13}$$

where y_0 is the observed values, $i = 1, \ldots, N_0$ from the observed dataset, and $\hat{\mathcal{L}}_h$ represents a penalty term corresponding to inequality constraints stated the third term in (2). The modifications from loss configurations of standard PINNs are the introducing multipliers λ for each categorized loss as loss terms and the weights m for each individual loss for the outputs on each state variable in the categorized loss.

3. Multi-Objective Optimization

The choice of weight coefficients for the different loss components requires careful tuning to balance the contributions of the residual loss, boundary/initial condition loss, and constraint-aware loss. As mentioned in (Lu et al., 2021), if the multiplier for the categorized loss is larger, the constraint violations are penalized more severely, forcing the solutions to better satisfy the constraints. However, when the penalty coefficients are too large, the optimization problem becomes ill-conditioned and difficult to converge to a minimum. On the other hand, if the penalty coefficients are too small, then the obtained solution will not satisfy the constraints and thus is not a valid solution. Although the soft-constraint approach has worked well for inverse problems to match observed measurements, it cannot be used in general because we cannot determine appropriate multipliers in the learning process.

To enhance the usability and robustness of the framework, automated techniques for optimal weight selection are employed. This study, involves a combination of two balanced processes in the learning process, inspired by (McClenny & Braga-Neto, 2020; Wang et al., 2023). The first balancing intensifies the gradient of individual losses in categorized losses to enhance local constraints, especially the objective for inequality derivative constraints. The second balancing addresses the multi-scale imbalance between categorized losses in (10∼13), particularly to mitigate the changing of gradient values from epoch to epoch due to the inequality feature in (13).

3.1. Individual Loss Balancing

In alignment with the neural network philosophy of self-adaptation, this study applies a straightforward procedure with fully trainable weights to generate multiplicative soft-weighting and attention mechanisms. The first balancing proposes self-adaptive weighting that updates the loss function weights via gradient ascent concurrently with the network weights. We minimize the total cost with respect to θ but also maximize it with respect to the self-adaptation weight vectors m at the k-th epoch,

$$m_\beta^{(j)}(k+1) = m_\beta^{(j)}(k) + \eta_m \nabla_{m_\beta^{(j)}} \hat{\mathcal{L}}_\beta(k). \tag{14}$$

where β specifies each loss $\beta \in \{0, b, f, h\}$ and η_m is learning parameter. In the learning step, the derivatives with respect to self-adaptation weights are increased when the constraints are violated and become larger when the errors are larger.

3.2. Categorized Loss Balancing

Parallely, for the second balancing, loss-balancing employs the following weighting function with balancing parameters λ in the loss function based on Eq (9). Considering updated at the k-th epoch,

$$\lambda_\beta(k+1) = \begin{cases} 1, & \text{if } \overline{\left|\nabla_\theta \hat{\mathcal{L}}_\beta(k)\right|} = 0 \\ \lambda_\beta(k) + \dfrac{\sum_\beta \overline{\left|\nabla_\theta \hat{\mathcal{L}}_\beta(k)\right|}}{\overline{\left|\nabla_\theta \hat{\mathcal{L}}_\beta(k)\right|}}, & \text{otherwise} \end{cases} \tag{15}$$

where ∇_i is the partial derivative vector (gradient) with respect to the i-th input vector (or value), and $\overline{|\cdot|}$ indicates the average of the absolute values of the elements in the vector. Note that the main reason for the choice of absolute values, instead of squared values as in (Wang et al., 2023), is to avoid overlooking outlier violations of the inequality constraints because most elements of $\nabla \mathcal{L}_h$ are assumed to be zero values in almost all cases.

The applied methods automatically adjust the weights of loss terms based on their relative magnitudes during training at user-specified intervals. These adaptive approaches have the potential to ensure a balanced contribution of each term of unstable inequality losses to the optimization process, potentially improving convergence and accuracy.

4. Neural Network Formulation

In this study, we apply a simple but deep feed-forward neural network architecture, a multilayer perceptron (MLP). Let $L \geq 2$ be an integer representing the depth of the network; we consider a neural network constructed with one input vector, L hidden layers, and one output value. The input values and the output variable are real numbers, i.e., $x \in \mathbb{R}^n$

and $y \in \mathbb{R}$. We can consider the price function φ in 2.1 which includes MLP as a multivariate function φ depending on the variables x, i.e., $\varphi : \mathbb{R}^2 \to \mathbb{R}$,

$$\varphi(x) = A_L \circ f_{L-1} \circ A_{L-1} \circ \cdots \circ f_1 \circ A_1(x), \quad (16)$$

where for $l = 1, \ldots, L$, $A_l : \mathbb{R}^{d_{l-1}} \to \mathbb{R}^{d_l}$ are affine functions as $A_l(x_{l-1}) = W_l^{\mathsf{T}} x_{l-1} + b_l$, and d_l is the number of neurons in the next layer l for $x_{l-1} \in \mathbb{R}^{d_{l-1}}$, with $W_l \in \mathbb{R}^{d_{l-1} \times d_l}$ and $b_l \in \mathbb{R}^{d_l}$, $d_0 = n$, $d_L = 1$, and $x_0 = x$. f_l are an activation function which is applied component-wise. Given a dataset x, which includes a set of pairs $(x^{(i)}, y^{(i)}), i = 1, \ldots N$, the network model φ is found by fitting the parameters θ (i.e., $W$and b) which minimize the designed loss function $\mathcal{L}$.

A challenge lies in that loss functions Eq. (9) for derivative-constrained PDEs involve derivatives with respect to input variables x, which are also functions of the parameters. When numerical approximation of derivatives is used, it could result in slow or inaccurate solutions. To address this, this study utilizes an extended backpropagation algorithm from (Rumelhart et al., 1986) with Automatic Differentiation for the gradient in Eq. (9) through exact derivative formulations, which require activation functions in the network to be second-order differentiable or higher. It is noted that some functions like Relu or Elu need slight additional consideration at non-differentiable singular points. If all activation functions are second-order differentiable, the same is true for the whole MLP, as shown in (Hornik et al., 1990).

5. Algorithms

This section introduces the algorithm of the DC-PINNs for multi-objective problems, which control the inequity loss of the partial derivatives of a neural network function with respect to its input features and apply the combination of loss balancing techniques for both categorized and individual losses.

Algorithm 1 exhibits DC-PINNs characteristics that set it apart from conventional learning methods. First, the computation points $x_{\{f,h\}}$ for the derivatives of the MLP do not correspond with the points of the training dataset x_0. The algorithm adjusts the derivatives to fit wide mesh grids in the defined space, thereby capturing derivative data across a wide array of input features. Secondly, the objective function $\mathcal{L}$ does not depend only on the MLP's direct output but also on its derivatives as specified in Eq.(9), all of which depend on identical network parameters. DC-PINNs facilitate balancing among categorized losses in addition to enhanced individual losses, which consist of PDE residuals and various scaled losses resulting from the violation of inequality constraints.

Algorithm 1 DC-PINNs with Balancing Processes

Input: Dataset $(x_{0,b}, y_{0,b})$, x_f, x_h, η, η_m, p_m, p_λ, $k_{\max}$

Consider a deep NN $\varphi_\theta(x)$ with θ, and a loss function

$$\mathcal{L} := \textstyle\sum_\beta \lambda_\beta \hat{\mathcal{L}}_\beta \left(m_\beta, x_\beta\left(, y_\beta\right)\right),$$

where $\hat{\mathcal{L}}_\beta$ denotes the categorized loss with $\beta \in \{0, b, f, h\}$, $m_\beta = \mathbf{1}$ are soft-weighting vectors for individual losses and $\lambda_\beta = 1$ are dynamic multipliers.

for $k = 1, \ldots, k_{\max}$ **do**

  Compute $\nabla_\theta \hat{\mathcal{L}}_\beta(k)$ by automatic differentiation

  if $k \equiv 0 \pmod{p_m}$ **then**

    Update m_β by

$$m_\beta^{(j)}(k+1) = m_\beta^{(j)}(k) + \eta_m \nabla_{m_\beta^{(j)}} \hat{\mathcal{L}}_\beta(k),$$

    where $m_\beta(k), \hat{\mathcal{L}}_\beta(k)$ shows values at k-th iteration.

  end if

  if $k \equiv 0 \pmod{p_\lambda}$ **then**

    Update λ_β by

$$\lambda_\beta(k+1) = \begin{cases} 1, & \text{if } \alpha = 0 \\ \lambda_\beta(k) + \frac{\sum_\beta \alpha}{\alpha}, & \text{otherwise} \end{cases}$$

    where $\alpha_\beta = \overline{\left|\nabla_\theta \hat{\mathcal{L}}_\beta(k)\right|}$

  end if

  Update the parameters θ via gradient descent, e.g.,

$$\theta(k+1) = \theta(k) - \eta \nabla_\theta \mathcal{L}(k).$$

end for

Return: θ

6. Experimental Design

6.1. Neural Network Setting and Training Configuration

In the experiments, the network architecture φ is a deep setting with four hidden layers ($L = 5$), each containing 32 neurons and a hyperbolic tangent activation function for smooth activations, following (Wang et al., 2023) for specific measures to improve learning efficiency and accuracy in selecting appropriate architectures. We also employ Glorot initialization for network parameters and the Adam optimization (Kingma & Ba, 2014) with a weight decay setting, which starts with a learning rate $\eta = 10^{-3}$ and an exponential decay with a decay rate of 0.9 for every 1000 decay steps. The hyperparameters used are $p_m, p_\lambda = 100$, and $k_{\max} = 10000$. The compared models in the result section are MLP $\mathcal{L} := \mathcal{L}_0 + \mathcal{L}_b$, PINNs $\mathcal{L} := \mathcal{L}_0 + \mathcal{L}_b + \mathcal{L}_f$, and DC-PINNs $\mathcal{L} := \lambda_0 \hat{\mathcal{L}}_0 + \lambda_b \hat{\mathcal{L}}_b + \lambda_f \hat{\mathcal{L}}_f + \lambda_h \hat{\mathcal{L}}_h$. Training data is prepared using equally distributed points for initial and boundary conditions ($N_0 = N_b = 101$) and containing square mesh grids for PDE residuals and inequality constraints, i.e., $N_f = N_h = 101 \times 101$. To evaluate approximation ability, the evaluation errors between predictions and answers are calculated using mesh grids as the same grids of constraints.

In computing, differentiable operators have been developed

in JAX/Flax (Bradbury et al., 2018; Heek et al., 2023), which can efficiently calculate exact derivatives using automatic differentiation. The experiments are conducted using Google Colab[1], which offers GPU computing on the NVIDIA Tesla T4 with a video random access memory of 15 GB.

7. Numerical Experiments

To demonstrate the effectiveness of the proposed framework, we conduct a series of numerical experiments on several benchmark problems related to quantitative finance: heat diffusion, Black-Scholes option pricing, and local volatility surface calibration with complex geometries and nonlinear constraints. We compare our approach to existing PINN-based approaches, including derivative profile comparisons on a function of trained networks.

7.1. One-dimensional Heat Equation in Thermodynamics

The heat equation is a classic example of a parabolic partial PDE problem (Cannon, 1984). It is a suitable and well-established problem for illustrating the robustness of PINNs. Consider an infinitesimally thin steel beam heated at its centre by a heat source. The heat at the centre will spread over the steel beam while the edges are kept at zero temperature, ensuring that the temperature reaches zero at an infinite final time. The problem setup is as follows:

$$f(x,t) = \frac{\partial y}{\partial t} - \lambda \frac{\partial^2 y}{\partial x^2}, \tag{17}$$

$$\text{s.t.} \quad y(x,0) = \sin(\pi x),\ y(0,t) = y(1,t) = 0, \quad \frac{\partial^2 y}{\partial x^2} \leq 0, \frac{\partial y}{\partial t} \leq 0, \tag{18}$$

where $x, t \in [0,1]$. The solution to this heat equation is given by $y(t,x) = e^{-\lambda \pi^2 t} \sin \pi x$. We can add derivative constraints, as shown in the second line of (18), based on the fact that the specified derivatives of the analytical solution should be negative in the defined space. This is also intuitively convincing because the high heat at the centre gradually spreads to the edges, maintaining solutions as parabolic curves over the domain throughout the entire timespan with heat reduction. We set coefficient $\lambda = 0.1$ in this experiment.

Figure 1 compares the temperature fields learned by the proposed DC-PINNs framework with those of the standard PINNs approach. Both methods appear to achieve sufficient fitting and do not exhibit significant differences in accuracy when fitting the overall temperature profile.

To further investigate the impact of incorporating inequality constraints in DC-PINNs, Figure 2 illustrates the derivative profiles of the learned temperature fields with respect to the spatial coordinate x at different time snapshots. These sensitivity profiles highlight the key differences between the PINNs and DC-PINNs solutions. The standard PINNs generate temperature profiles with distinct regions of non-physical positive in differentials $\frac{\partial^2 y}{\partial x^2}$ and $\frac{\partial y}{\partial t}$. In contrast, DC-PINNs consistently produce sensitivity profiles that adhere to non-positivity constraints in differentials $\frac{\partial^2 y}{\partial x^2}$ and $\frac{\partial y}{\partial t}$ across all time snapshots while closely matching the ground truth profiles. This demonstrates the effectiveness of the DC-PINNs framework in enforcing inequality constraints during the learning process, resulting in physically consistent solutions. These results emphasize the importance of explicit constraint handling in physics-informed neural networks, as naively relying on the partial differential equation (PDE) residual term alone may lead to solutions that violate critical physical principles.

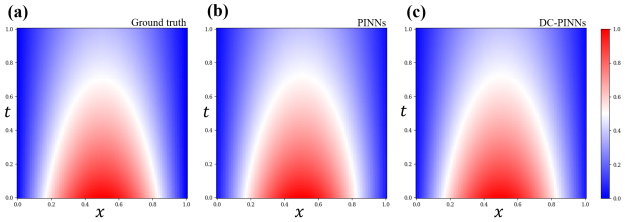

Figure 1. Numerical solutions for the one-dimensional heat equations with derivative-constrained conditions. (a) Exact solution. (b) PINNs solution. (c) DC-PINNs solution.

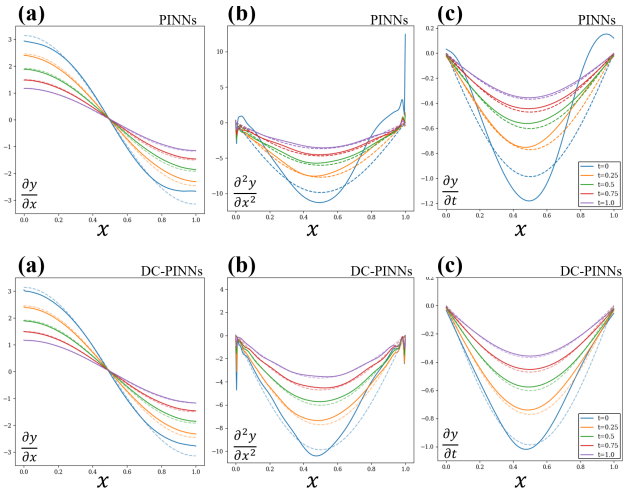

Figure 2. Derivative profiles of the learned temperature fields at different time snapshots for PINNs (upper row) and DC-PINNs (bottom row), aligned with exact solutions (dashed lines). (a) First-order derivatives with respect to x. (b) Second-order derivatives with respect to x. (c) First-order derivatives with respect to t.

[1]Google Colab. http://colab.research.google.com

7.2. Black-Scholes Model in Quantitative Finance

Next, as an example of solving a PDE in which inequality constraints on derivatives play a major role, we consider the European (call) option pricing problem, which is a typical quantitative finance problem governed by the Black-Scholes model (Black & Scholes, 1973). The Black-Scholes equation is a parabolic partial differential equation that describes the option price y in terms of the price of the underlying asset x and time t, given the following initial and boundary conditions:

$$f(x,t) = \frac{\partial y}{\partial t} + \frac{1}{2}\sigma^2 x^2 \frac{\partial^2 y}{\partial x^2} + rx\frac{\partial y}{\partial x} - ry, \tag{19}$$

$$\begin{aligned} \text{s.t.} \quad & y(x,t_m) = (x_{t_m} - k)^+, \; y(0,t) = 0, \\ & \frac{\partial y}{\partial x} \geq 0, \; \frac{\partial^2 y}{\partial x^2} \geq 0, \; \frac{\partial y}{\partial t} \leq 0, \end{aligned} \tag{20}$$

where k is the strike price, $\tau = t_m - t$ is the time to maturity t_m, σ is the volatility, and r is the risk-free rate. For the purpose of this study, we consider a specific type of European option with $x, t \in [0,1]$, $k = 0.5$, $t_m = 1$, $r = 0.1$, and $\sigma = 0.3$. The inequalities in the second line of Equation (20) represent the necessary and sufficient conditions for no-arbitrage, which ensure that the option price is consistent with other financial strategies. These conditions are fundamental principles that the price of the financial product must satisfy. The no-arbitrage constraints are expressed as inequalities for the first and second derivatives with respect to the underlying asset x and time t, as discussed in Section A.2. The exact solutions for the option price with respect to x and t are provided by the Black-Scholes formula (Black & Scholes, 1973),

$$\begin{aligned} y(x,t) &= x_t N(d_+) - e^{-r\tau} k N(d_-), \\ d_\pm &= \frac{\ln(e^{-r\tau} x_t/k) \pm (\sigma^2/2)\tau}{\sigma\sqrt{\tau}}, \end{aligned} \tag{21}$$

where $N(\cdot)$ is the cumulative normal distribution function.

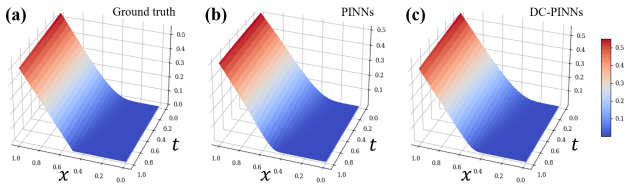

Figure 3. Numerical solutions for the Black-Scholes model in finance with derivative-constrained conditions. (a) Exact solution. (b) PINNs solution. (c) DC-PINNs solution.

Figure 3 illustrates the numerical solutions obtained by each method in a defined space. Both methods in the Black-Scholes pricing problem also appear to achieve appropriate fitting and do not exhibit significant differences in accuracy when fitting the overall price profile.

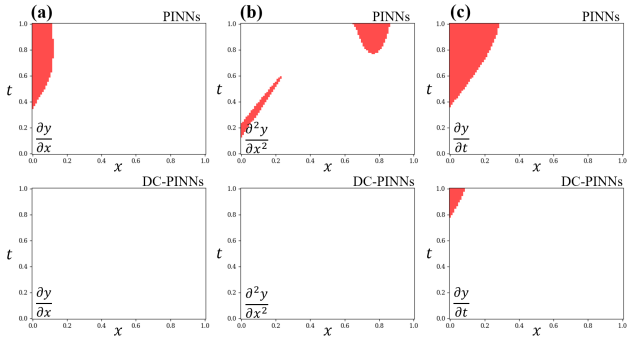

Figure 4. Heatmaps of the inequality derivatives conditions of the Black-Scholes option pricing for PINNs (upper row) and DC-PINNs (bottom row), where the area that violates constraints is coloured red. (a) First-order derivatives with respect to x. (b) Second-order derivatives with respect to x. (c) First-order derivatives with respect to t.

To further investigate the impact of incorporating inequality constraints in DC-PINNs as previous experiment, as in the previous example, Figure 4 illustrates heatmaps of the inequality derivatives conditions of the trained models; the area that violates constraints are coloured red. These sensitivity profiles also highlight the differences in solutions, which is that DC-PINNs consistently produce sensitivity profiles that adhere to the no-arbitrage constraints across almost time snapshots, although PINNs violate the inequality condition of derivatives on the wide area.

7.3. Implied Volatility Surface Calibration in Finance

Finally, we consider the calibration problem, an inverse problem to identify governing parameters in a PDE, where the option prices satisfy the PDE (i.e., the Local Volatility (LV) model explained in A). The implied volatility y is given with respect to strike x and time to maturity t by,

$$\begin{aligned} f(x,t) &= \frac{\partial \tilde{y}}{\partial t} - \frac{1}{2}\sigma_{\text{LV}}^2 x^2 \frac{\partial^2 \tilde{y}}{\partial x^2} + rx\frac{\partial \tilde{y}}{\partial x}, \\ \tilde{y} &= g(x,t,y), \quad \sigma_{\text{LV}} = \phi(x,t,y,\mathcal{D}y), \end{aligned} \tag{22}$$

$$\begin{aligned} \text{s.t.} \quad & \tilde{y}_{t=0} = (s_{t_m} - x)^+, \; \tilde{y}_{x=0} = s_t, \\ & \frac{\partial \tilde{y}}{\partial x} \leq 0, \; \frac{\partial^2 \tilde{y}}{\partial x^2} \geq 0, \; \frac{\partial \tilde{y}}{\partial t} \geq 0, \end{aligned} \tag{23}$$

where $\tilde{y}$ represents the option price calculated by the function $g(\cdot)$, which is plugged into the formula in (25). The conversion function $\phi(\cdot)$ is with respect to y using (28). Assuming European (call) options with various strikes and time to maturity $x, t \in [0,1]$, $s_0 = 0.1$, $t_m = 1$, $r = 0.1$, and $s_t = s_0 e^{rt}$. Synthetic data is prepared as option premiums

($\tilde{y}$) using the SABR model, as described in A.4. It is noted that this problem is different from previous examples since the feasible solution by optimizations affects the designed parameters of PDEs.

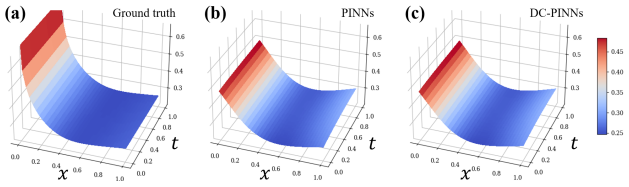

Figure 5. Numerical solutions (σ_{LV}) for the implied surface calibration in finance with derivative constrained conditions. (a) Exact solution. (b) PINNs solution. (d) DC-PINNs solution.

Figure 5 illustrates the numerical solutions obtained by each method. Similar to the previous examples, both methods in the calibration problem also appear to achieve appropriate fitting and do not exhibit significant differences in accuracy when predicting the overall parameter profiles.

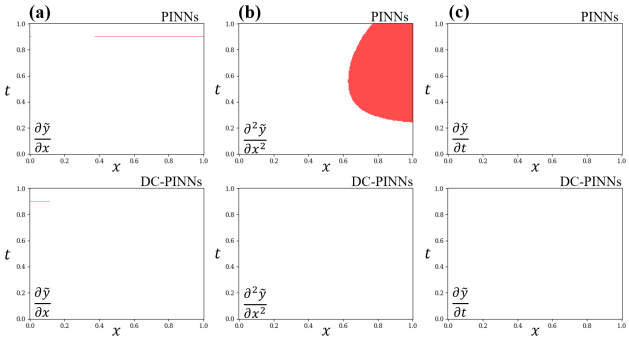

Figure 6. Heatmaps of the inequality derivatives conditions of the learned local volatilities for PINNs (upper row) and DC-PINNs (bottom row), the area that violates constraints is coloured red. (a) First-order derivatives with respect to x. (b) Second-order derivatives with respect to x. (c) First-order derivatives with respect to t.

As with other experiments, we compare the performance of DC-PINNs that incorporate nonlinear constraints by investigating the derivative profiles of the obtained surface. Figure 6 illustrates the area of heatmaps of the inequality derivatives conditions of the learned local volatilities as same as previous examples. The results demonstrate that the proposed DC-PINNs framework captures the solution while satisfying the nonlinear constraints, outperforming the traditional PINNs approach. Furthermore, the DC-PINNs framework demonstrates versatility, as it can be applied to calibration problems of design parameters in PDEs, including data-driven solutions.

7.4. Computational Efficiency

Table 1. The computation times (in seconds) for the training in 7.3 based on changes in dataset size (total N) and number of neurons.

Models	Default	Dataset size		♯ of neurons	
		Half	Quarter	Half	Quarter
MLP	48.2	45.0	44.8	45.4	44.8
PINNs	95.5	68.4	56.7	70.6	57.9
DC-PINNs	122.3	92.6	80.9	93.2	82.6

At last, we demonstrate the effectiveness of DC-PINNS during training and evaluate its computational efficiency. Table 1 presents the computation times required by DC-PINNs and the baseline models for calibrating the surface in 7.3 based on changes in dataset size and number of neurons. The computational efficiency of DC-PINNs is evident from their ability to handle these complex optimization challenges without significant overhead. The efficient use of automatic differentiation and the adaptive loss balancing approach contribute to DC-PINNs convergence and reduced computational overhead. DC-PINNs also exhibit reasonable scalability, as evidenced by the sublinear growth in computation time with respect to the dataset size. This scalability is crucial for handling the ever-increasing volumes of real-world data in modern scientific domains.

8. Conclusion and Future Work

In this study, we proposed an extended PINNs framework called Derivative-Constrained PINNs (DC-PINNs) to effectively solve PDEs with derivative constraints, which seamlessly integrates constraint information into the learning process through a constraint-aware loss function. The effectiveness of the DC-PINNs framework was demonstrated through a series of numerical experiments on benchmark problems related to quantitative finance, including heat diffusion, Black-Scholes pricing, and local volatility surface calibration. The results showed that DC-PINNs outperformed standard PINNs approaches in terms of the ability to satisfy nonlinear constraints with sufficient accuracy. DC-PINNs also illustrated computational efficiency in handling optimization without significant overhead by the use of automatic differentiation and adaptive loss balancing techniques.

The study emphasizes the importance of explicit constraint handling in PINNs, as relying solely on the PDE residual term may lead to solutions that violate critical physical principles. The DC-PINNs framework opens up new possibilities for solving constrained PDEs in multi-objective optimization problems. However, further investigation is needed to assess the scalability of the framework to high-dimensional and more complex constraint types. Future work could explore efficient sampling strategies and adap-

tive collocation point selection to mitigate the curse of dimensionality and improve accuracy and computational efficiency.

References

Black, F. The pricing of commodity contracts. *Journal of financial economics*, 3(1-2):167–179, 1976.

Black, F. and Scholes, M. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3): 637–654, 1973.

Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/google/jax.

Cannon, J. R. *The one-dimensional heat equation*. Number 23. Cambridge University Press, 1984.

Carr, P. and Madan, D. B. A note on sufficient conditions for no arbitrage. *Finance Research Letters*, 2(3):125–130, 2005.

Chen, Y., Huang, D., Zhang, D., Zeng, J., Wang, N., Zhang, H., and Yan, J. Theory-guided hard constraint projection (hcp): A knowledge-based data-driven scientific machine learning method. *Journal of Computational Physics*, 445: 110624, 2021.

Dupire, B. et al. Pricing with a smile. *Risk*, 7(1):18–20, 1994.

Hagan, P. S., Kumar, D., Lesniewski, A. S., and Woodward, D. E. Managing smile risk. *The Best of Wilmott*, 1: 249–296, 2002.

Heek, J., Levskaya, A., Oliver, A., Ritter, M., Rondepierre, B., Steiner, A., and van Zee, M. Flax: A neural network library and ecosystem for JAX, 2023. URL http://github.com/google/flax.

Hornik, K., Stinchcombe, M., and White, H. Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural networks*, 3(5):551–560, 1990.

Jagtap, A. D., Kharazmi, E., and Karniadakis, G. E. Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 365:113028, 2020.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Lagaris, I. E., Likas, A., and Fotiadis, D. I. Artificial neural networks for solving ordinary and partial differential equations. *IEEE transactions on neural networks*, 9(5): 987–1000, 1998.

Lo, K. and Huang, D. On training derivative-constrained neural networks. *arXiv preprint arXiv:2310.01649*, 2023.

Lu, L., Pestourie, R., Yao, W., Wang, Z., Verdugo, F., and Johnson, S. G. Physics-informed neural networks with hard constraints for inverse design. *SIAM Journal on Scientific Computing*, 43(6):B1105–B1132, 2021.

McClenny, L. and Braga-Neto, U. Self-adaptive physics-informed neural networks using a soft attention mechanism. *arXiv preprint arXiv:2009.04544*, 2020.

Raissi, M., Perdikaris, P., and Karniadakis, G. E. 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.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.

Wang, S., Sankaran, S., Wang, H., and Perdikaris, P. An expert's guide to training physics-informed neural networks. *arXiv preprint arXiv:2308.08468*, 2023.

A. Implied Volatility Surface Calibration

A.1. Calibration Problem

The implied volatility surface is a model of values resulting from European option prices. We are given a complete filtered probability space $\left(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P}\right)$ and that $\mathbb{P}$ is an associated risk-neutral measure. The price of a European call option C at time t is defined as

$$C = e^{-r\tau}\mathbb{E}\left[(S_T - K)^+ \mid \mathcal{F}_t\right], \tag{24}$$

where S_t is the underlying price at t, K is the strike price, $\tau = T - t$ is the time to maturity T, and r is the risk-free rate. In (Black & Scholes, 1973), the implied volatility σ_{imp} leads to the modelled price with $K, \tau \in [0, \infty)$ as the Black-Scholes (BS) formula,

$$\begin{aligned} C_{\text{BS}}(\sigma_{\text{imp}}) &= S_t N(d_+) - e^{-r\tau} K N(d_-), \\ d_\pm &= \frac{\ln(e^{-r\tau} S_t/K) \pm \left(\sigma_{\text{imp}}^2/2\right)\tau}{\sigma_{\text{imp}}\sqrt{\tau}}, \end{aligned} \tag{25}$$

where $N(\cdot)$ is the cumulative normal distribution function.

To expand the generalization for K and τ, (Dupire et al., 1994) proposed the Local Volatility (LV) model, in which the European option prices satisfy the PDE,

$$rK\frac{\partial C}{\partial K} - \frac{1}{2}\sigma_{\text{LV}}^2(K,\tau)K^2\frac{\partial^2 C}{\partial K^2} + \frac{\partial C}{\partial \tau} = 0, \tag{26}$$

with initial and boundary conditions given by

$$C_{\tau=0} = (S_T - K)^+, \quad \lim_{K\to\infty} C = 0, \quad \lim_{K\to 0} C = S_t. \tag{27}$$

Plugging it into the formula in (25), one obtains a conversion function σ_{imp} and σ_{LV} with respect to K and τ,

$$\sigma_{\text{LV}}^2(K,\tau) = \frac{\sigma_{\text{imp}}^2 + 2\sigma_{\text{imp}}\tau\left(\frac{\partial\sigma_{\text{imp}}}{\partial\tau} + rK\frac{\partial\sigma_{\text{imp}}}{\partial K}\right)}{1 + 2d_+K\sqrt{\tau}\frac{\partial\sigma_{\text{imp}}}{\partial K} + K^2\tau\left(d_+d_-\left(\frac{\partial\sigma_{\text{imp}}}{\partial K}\right)^2 + \sigma_{\text{imp}}\frac{\partial^2\sigma_{\text{imp}}}{\partial K^2}\right)} \tag{28}$$

to be fit with the PDE (26).

We can define the calibration as identifying the multivariate function respecting the prices $\Phi(x)$, associated with implied volatility surface as a function $\varphi(x) \geq 0$ with inputs $x := (K, \tau)$,

$$\Phi(x) = C_{\text{BS}}(K, \tau, \varphi(x)). \tag{29}$$

The inverse problem of the implied volatility surface is that, given limited options prices, we would like to identify the implied volatility function with respect to x, redefined as $\sigma_{\text{imp}}(x)$, to fit with that premium resulted by C_{BS} also satisfy PDE. Based on (25∼28), once φ is determined, we can analytically obtain the option price Φ. Furthermore, Φ is second differentiable whenever φ is second differentiable, allowing the representation of the PDE in (26).

A.2. No-Arbitrage Constraints for European Options

The option prices should obey the constraints imposed by no-arbitrage conditions, which are essential financial principles posit that market prices prevent guaranteed returns above the risk-free rate. This study considers the necessary and sufficient conditions for no-arbitrage presented in (Carr & Madan, 2005). This allows us to express the call option price as a two-dimensional surface appropriately. The necessary and sufficient conditions for no-arbitrage are represented as non-strict inequalities for several first and second derivatives,

$$-e^{-r\tau} \leq \frac{\partial C}{\partial K} \leq 0, \quad \frac{\partial^2 C}{\partial K^2} \geq 0, \quad \frac{\partial C}{\partial \tau} \geq 0. \tag{30}$$

From the above, no-arbitrage conditions require these derivatives to have a specific sign. The standard architecture does not automatically satisfy these conditions when calibrating with a loss function simply based on the mean squared error (MSE) for the prices.

A.3. The SABR model

The SABR model in (Hagan et al., 2002) is a typical parametric model, which can capture the market volatility smile and skewness and reasonably depict market structure. When F_t is defined as the forward price of an underlying asset at time t, the SABR model is described as

$$\begin{aligned} dF_t = \alpha_t F_t^{\beta} dW_t^1, \quad d\alpha_t = \nu \alpha_t \, dW_t^2, \\ \langle dW_t^1, dW_t^2 \rangle = \rho dt. \end{aligned} \tag{31}$$

Here, W_t^1, W_t^2 are standard Wiener processes, α_t is the model volatility, ρ is the correlation between the two processes, and ν is analogous to vol of vol. The additional parameter β describes the slope of the skewness. Essentially, the IV in the SABR model is given by a series expansion technique associated with volatility form of (Black, 1976)

$$\begin{aligned} \sigma(K,\tau) &= \frac{\alpha \left(1 + \left(\frac{(1-\beta)^2}{24} \frac{\alpha^2}{(FK)^{1-\beta}} + \frac{1}{4} \frac{\rho\beta\upsilon\alpha}{(FK)^{(1-\beta)/2}} + \frac{2-3\rho^2}{24} \upsilon^2\right) \tau\right)}{(FK)^{(1-\beta)/2} \left[1 + \frac{(1-\beta)^2}{24} \ln^2 \frac{F}{K} + \frac{(1-\beta)^4}{1920} \ln^4 \frac{F}{K}\right]} \frac{z}{\chi(z)}, \\ z &= \frac{\upsilon}{\alpha} (FK)^{(1-\beta)/2} \ln \frac{F}{K}, \quad \chi(z) = \ln \left(\frac{\sqrt{1 - 2\rho z + z^2} + z - \rho}{1 - \rho} \right). \end{aligned} \tag{32}$$

A.4. Testing with Synthetic Data

The algorithm developed for evaluating volatility surface was first tested on simulated values in a parameterized two-dimensional case of the surface interpolation problem. We took up the Stochastic Alpha Beta Rho (SABR) model introduced by (Hagan et al., 2002) and prepared a sparse two-dimensional dataset to test our methodology. To fit a more realistic market situation, the following experiment utilized sparse (and uneven) grid data referring to $10,000$ random grids $x, t \in [0, 1]$ and generates synthetic option premiums on the grids by Eq (32) with $\{\alpha, \beta, \rho, \nu\} = \{0.3, 1.0, -0.6, 0.4\}$.