NEURAL SUPERPOSITION NETWORKS

Anonymous authors

Paper under double-blind review

Abstract

Machine learning models can be biased towards the solutions of given differential equations in two principal ways: through regularisation, or through architecture design. Recent research has successfully constrained neural network architectures to satisfy divergence-free fields and Laplace's equation in two dimensions. This work reinterprets these architectures as linear superpositions of general formulated solutions. The notion of superposition is then exploited to develop novel architectures for Laplace's equation and divergence-free fields, we propose novel constraints apt for the heat equation, and even some nonlinear differential equations including Burgers' equation. Benchmarks of superposition-based approaches against previously published architectures and physics-informed regularisation approaches are presented. We find that embedding differential equation constraints directly into neural network architectures can lead to improved performance and hope our results motivate further development of neural networks architectures developed to adhere specifically to given differential constraints.

023

000

001 002 003

004

006 007

008 009

010

011

012

013

014

015

016

017

018

019

021

1 INTRODUCTION

025 026

027 There has been fruitful interplay between the fields of machine learning and differential equation solving. Differential equations provided useful means to machine learning: ordinary differential 029 equations (ODEs) appear in continuous depth neural networks (Chen et al., 2018) and stochastic differential equations (SDEs) inspired novel approaches to generative modelling (Song et al., 2020). On the other side, neural network architectures have also provided novel means to analyse and solve 031 differential equations, in particular, physics-informed neural networks (PINNs) (Lagaris et al., 1998; Raissi et al., 2019) bias neural network outputs towards a given differential equation via construction 033 of appropriate regularisation terms. Also, neural networks have been shown to be effective surrogate 034 models of computationally-expensive, finite element solvers of partial differential equations (PDE) (Nabian & Meidani, 2018). In addition to providing novel means for solving forward problems of differential equations, neural networks proved useful in solving inverse-problems in differential 037 equations (Lu et al., 2021b), data-driven discovery of differential equations (Both et al., 2021) and 038 approaches for optimal control (Mowlavi & Nabi, 2023).

Thus far, physics-informed regularisation has been a dominant approach to biasing machine learning models towards differential equations, being widely applied from PDEs and ODEs (Raissi et al., 2019) to fractional order differential equations (Pang et al., 2019). Despite their widespread applicability, the difficulty of training PINNs has been highlighted in many works that also highlight potential means to improve their training. For example, via appropriate loss-scaling techniques (Wang et al., 2021; Son et al., 2023), domain decomposition (Jagtap & Karniadakis, 2020; Moseley et al., 2023), combining gradient and Hessian-based optimisation procedures (Mao et al., 2020), and relaxation of the underlying differential equation form (Krishnapriyan et al., 2021).

It is yet unclear how to reliably train PINNs, or eventually provide convergence guarantees. Therefore, recent research has started to move beyond PINNs, towards embedding target differential equations into underlying architectures directly. For example, neural conservation laws (NCLs), which
involve post-processing of multilayer perceptrons (MLPs) with higher-order derivatives to yield
divergence-free neural networks, ensure conservation principles by using divergence-free vector
fields (Richter-Powell et al., 2022). Holomorphic neural networks, which are complex-valued MLPs
with appropriate activation functions, have been shown to satisfy the Laplace equation (Ghosh et al., 2023). Additionally, neural networks satisfying Hamiltonian laws have been developed to preserve

the total energy of a dynamical system over time (Greydanus et al., 2019). Beyond the context of neural networks, priors for Gaussian processes over the space of solutions of given linear differential equations have also been developed (Harkonen et al., 2023).

058 **Contributions of this work** This work presents a framework for constructing neural networks to satisfy linear differential equations through the principle of superposition. Since the solution space 060 of homogeneous linear differential equations is closed under addition, we can derive single-layer feedforward architectures that leverage a superposition of solutions as a learnable representation of 061 062 solution functions. We demonstrate how this framework provides a unifying perspective on previously published neural network architectures applying to divergence-free vector fields (Richter-063 Powell et al., 2022) and Laplace's equations (Ghosh et al., 2023). Novel architectures expanding on 064 such applications to encompass also the heat equation and Burgers' equation are presented. Despite 065 Burgers' equation being nonlinear, we are able to construct appropriate architectures constrained to 066 it via suitable post-processing of architectures satisfying the heat equation. We benchmark our archi-067 tectures against representative, standard architectures constrained to the same differential equations 068 (where applicable) as well as physics-informed neural networks. 069

2 Theory

We consider neural network architectures constrained to satisfy homogeneous linear differential equations.

Consider an open set $\Omega \subset \mathbb{R}^{d_1}$, $d_1 \in \mathbb{N}$ with boundary $\partial \Omega$. Define a space of sufficiently smooth functions, \mathcal{A} , from $\Omega \cup \partial \Omega$ to \mathbb{R}^{d_2} , $d_2 \in \mathbb{N}$. We consider equations of the form

077 078

071 072

057

079

093

094

095 096 097

100

 $\mathcal{L}u = 0 \text{ in } \Omega$ $\mathcal{N}u = 0 \text{ on } \partial\Omega$ (1)

where $u \in \mathcal{A}$, \mathcal{L} and \mathcal{N} are linear differential operators mapping elements of \mathcal{A} to \mathbb{R}^{d_3} -valued functions on $\Omega \cup \partial \Omega$, i.e. that $\mathcal{L}(af + bg) = a\mathcal{L}f + b\mathcal{L}g$ for all $f, g \in \mathcal{A}$ and $a, b \in \mathbb{R}$. In the case where $d_3 > 1$ is vector-valued, we take the right-hand side to be the zero-vector.

Note that the assumptions in our formulations encompass a wide-range of common differential equations with wide-reaching practical applications, such as the Laplace, Diffusion, Heat, Sturm-Liouville and Wave equations, to name but a few. In all these cases, the goal is to find a function *u* satisfying the differential equation and boundary conditions imposed by respectively \mathcal{L} and \mathcal{N} in Eq. equation 1.

Define a family of parameterised functions, $u_{\theta_i}^i$, i = 1, ..., n for $n \in \mathbb{N}$, with $u_{\theta_i}^i : \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$ and θ_i representing trainable parameters. Note that superscripts denote positional elements of u; they do not represent powers. Interpreting each $u_{\theta_i}^i$ to be a row vector comprising of d_2 columns, we can represent the family of functions $u_{\theta_i}^i$ as an $n \times d_2$ matrix:

 $\mathbf{u} = \begin{pmatrix} u_{\theta_1}^{1(1)} & \dots & u_{\theta_1}^{1(d_2)} \\ \vdots & \ddots & \vdots \\ u_{\theta_n}^{n(1)} & \dots & u_{\theta_n}^{n(d_2)} \end{pmatrix}$ (2)

we allow u to be interpreted as a matrix-valued function operating on $\Omega \cup \partial \Omega$, where the u(x) is determined by the element-wise application of each element of u to x.

Definition 1 (Superposition network) Given an $1 \times n$ matrix W, a function $\phi : (\Omega \cup \partial \Omega) \to \mathbb{R}^{d_2}$, given by $\phi_{\theta} = W\mathbf{u}$, where $\theta = \{W, \theta_1, \theta_2, \dots, \theta_n\}$ is referred to as a superposition network with respect to \mathcal{L} if and only if $\mathcal{L}u^i_{\theta_i} = 0$ for all i.

Given a superposition network with respect to \mathcal{L} denoted by ϕ_{θ} , it follows directly that $\mathcal{L}\phi_{\theta} = 0$ by linearity of \mathcal{L} . Thus, by solving the following optimisation procedure:

107

$$\theta^* = \operatorname*{arg\,min}_{\theta} \mathbb{E}_{x \sim \mathbb{P}_{\partial \Omega}} \left[\left(\mathcal{N} \phi_{\theta}(x) \right)^2 \right], \tag{3}$$

where the expectation is taken with respect to a probability distribution $\mathbb{P}_{\partial\Omega}$ defined over $\partial\Omega$. The resulting function ϕ_{θ^*} then approximates the solution of Eq. equation 1, as exemplified in Fig. 1.

The construction of superposition networks with respect to a given linear-differential operator \mathcal{L} depends on the specific operator at hand. This is exemplified in section 2.1, where we note that existing architectures in the literature constrained to Laplace's equation and divergence-free constraints can be interpreted as superposition networks, and section 2.2 where we propose new architectures for additional operators.



127 128 Figure 1:

Figure 1: A schematic of superposition networks, a single-layer feedforward neural network architecture constrained to be in the solution space of a linear differential equation. Superposition networks use a library of known solutions of the differential equation (a) and apply Lie group symmetries derived from the differential equation to derive suitable linear transformations (b) which can then be combined in an output layer (c) to approximate nontrivial solutions of the differential equation by training only on initial and boundary conditions.

133 134 135

136

147

148

149

150

160

129

130

131

132

2.1 EXISTING SUPERPOSITION NETWORKS

This section demonstrates that the principle of superposition provides a unifying perspective on previous neural network architectures used to model divergence-free fields (Richter-Powell et al., 2022) ($\mathcal{L} = \nabla \cdot$) and Laplace's equation (Ghosh et al., 2023) ($\mathcal{L} = \nabla^2$). Both of these works apply a postprocessing step to an existing neural network to constrain it to a certain differential equation. To unify our treatment of both architectures we introduce the notion of a linear postprocessing operator.

142 143 **Definition 2 (Linear postprocessing operator)** Consider a function $f : \mathbb{R}^{d_1} \to \mathbb{R}^{d_3}$, with com-144 ponents enumerated by $[f_1, f_2, \ldots, f_{d_3}]$. We equate f with its vector representation. Define an 145 operator A acting on f as a linear postprocessing operator with respect to \mathcal{L} if and only if the 146 following hold:

- 1. The action of A on f is \mathbb{R} -linear in the components of f, i.e. $A[af_1, af_2, \ldots, af_{d_3}] = aA[f_1, f_2, \ldots, f_{d_3}] \forall a \in \mathbb{R}$ and
- 2. $\mathcal{L}(A(f)) = 0.$

We use the notation $A[f_1, \ldots, f_{d_3}]$ to denote the application of A to f, which we write in a vector notation above to make clear that linearity of the operator acts on output components of f as opposed to the inputs of f.

Example 1 (Divergence-free neural networks) In work on modelling conservation laws, neural models of divergence-free vector-fields were derived by considering the Hodge-star operation on the exterior derivative of a differential form (Richter-Powell et al., 2022). The composition of the Hodge-star operator on the exterior derivative is a linear postprocessing operator with respect to the divergence operator.

Example 2 (Holomorphic neural networks) Taking the real part of a complex-valued MLP with holomorphic activation functions yields a neural network architecture which satisfies Laplace's

162 163 164 162 164 162 164 equation, $\nabla^2 \phi = 0$ (Ghosh et al., 2023). In this scenario, taking the real-part of the complex output of the MLP forms a linear postprocessing operator with respect to the Laplacian.

165 If f is taken as an MLP with a linear output layer, we can write $f = [f_1, f_2, ..., f_{d_3}]$ and take each 166 $f_i(x) = \sum_j w_{ij}\sigma_j(x)$, where $\sigma_j(x)$ represents the activation of the j-th neuron of the last hidden 167 layer of the MLP with input x. Application of the linear postprocessing operator can now yield a 168 degeneracy of ways of representing a target neural network. For example, applying linearity with 169 i = 1 yields:

$$A[\sum_{j} w_{1j}\sigma_{1}, \dots, f_{d_{3}}] = \sum_{j} |w_{1j}| A[\frac{w_{1j}}{|w_{1j}|} |\sigma_{1}, \dots, \frac{1}{|w_{1j}|} f_{d_{3}}],$$
(4)

with applications of linearity to other indices giving alternative superposition representations.

Note that each w_{ij} can potentially be complex-valued, as in the case for holomorphic neural networks (Ghosh et al., 2023), whereas the definition of linear postprocessing operators only assumes linearity over real numbers. Consequently, we only apply linearity to the absolute value of w_{1j} in Eq. equation 4.

The right hand side of Eq. equation 4 yields the definition of a superposition network if $\mathcal{L}A[\frac{w_{1j}}{|w_{1j}|}|\sigma_1,\ldots,\frac{1}{|w_{1j}|}f_{d3}] = 0$ for each j. For instance, this holds for the neural networks presented in Examples 1-2, which can thus be interpreted as special cases of superposition networks.

181 182

183

170 171

2.2 NOVEL SUPERPOSITION NETWORKS

The use of linear postprocessing operators in previous works (Ghosh et al., 2023; Richter-Powell et al., 2022) is tailored to specific linear differential equations—they use specific characteristics of 185 the exterior calculus of differential forms, or of holomorphic complex-valued functions, to cater to 186 particular differential equations. These approaches are not applicable to arbitrary linear differential 187 equations that we might consider as in Eq. equation 1. This section outlines an applicable approach 188 to systematically address new differential equations. We leverage the Lie group symmetries admitted 189 by a given differential operator to construct superposition networks applicable to cases beyond the 190 current state-of-art. Means to handle some non-linear differential equations are presented (i.e. the 191 Burgers' equation) in section 2.2.4.

Consider superposition networks of the form $\phi : \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$, given by:

196

$$\phi_{\theta}(x) = \theta_0 + \frac{1}{N} \sum_{i=1}^{N} \theta_1^{(i)} \sigma_i(g_{\theta_2^{(i)}}^{\sigma_i} x)$$
(5)

where the output layer of the architecture is explicitly summed and the following terms are introduced: (1) The activation function for the *i*th neuron, $\sigma_i : \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$, satisfies $\mathcal{L}\sigma_i = 0$ for all *i*. These activation functions can be trivial manufactured solutions for boundary conditions other than those specified in Eq. equation 1. (2) Linear transformations which are not freely chosen, but rather elements of a Lie group $g_{\theta_2^{(i)}}^{\sigma_i}$ acting on \mathbb{R}^{d_1} parameterised by trainable parameters $\theta_2^{(i)}$ such that $\mathcal{L}\sigma_i(g_{\theta_2^{(i)}}^{\sigma_i}x) = 0$ whenever $\mathcal{L}\sigma_i(x) = 0$. (3) Parameters associated with an affine output layer as usual: θ_0 and $\theta_1^{(i)}$.

Note that the inclusion of the parameter $\theta_0 \in \mathbb{R}^{d_2}$ generalises definition 1. However, in all differential equations we consider in this work, \mathcal{L} has first-order derivatives which allows us the flexibility to include θ_0 in our architecture. Taking $\theta = \left\{ W, \theta_0, \theta_2^{(1)}, \theta_2^{(2)}, \dots, \theta_2^{(N)} \right\}$, with $W = \left\{ \theta_1^{(1)}, \theta_1^{(2)}, \dots, \theta_1^{(N)}, \right\}$ allows us to the approximate a solution to Eq. equation 1 by solving the optimisation problem in Eq. equation 3.

Therefore, constructing appropriate superposition networks for a given linear-differential operator amounts to choosing suitable forms of σ_i and $g_{\theta_2^{(i)}}^{\sigma_i}$. In general, this can be done in a methodical way by enumerating the Lie group symmetries of \mathcal{L} (Gray, 2015). However, as we now demonstrate, for several important classes of practical equations, suitable forms can be constructed by observation and the method of manufactured solutions.

216 2.2.1 LAPLACE'S EQUATION

224

259 260

264 265

266

In two-dimensions, Laplace's equation is defined by taking $\mathcal{L} = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ in Eq. equation 1. While architectures have previously been designed to satisfy the Laplace equation, we present alternative architectures here as a means of exemplifying the construction of superposition networks.

It follows directly that the following two-parameter Lie groups translate solutions of the Laplace equation (harmonic functions) to other solutions of the Laplace equation:

$$(\hat{x}, \hat{y}) = g^D_{\theta_D} g^{E2}_{\theta_E}(x, y) \tag{6}$$

with (trainable) Lie group parameters being given by $\theta_E \in \mathbb{R}^3$ and $\theta_D \in \mathbb{R}$. The transformations in Eq. equation 6 are a composition of E2, the Euclidean group in two-dimensions, with a dilation of the two-dimensional plane given by $g^D_{\theta_D}(x, y) = (\theta_D x, \theta_D y)$.

For any function, $\sigma(x, y)$ such that $\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right)\sigma(x, y) = 0 \ \forall (x, y) \in \Omega$, it is readily verifiable that the Laplacian is invariant under group actions as described in Eq. equation 6, which thus define a suitable form for use as the transformation group in Eq. equation 5.

To define suitable functions σ_i in Eq. equation 5, we can take the real part of any holomorphic function $f : \mathbb{C} \to \mathbb{C}$, a choice for f(z) = f(x + iy) of for example $\sin z$, $\sin(\sin z)$, $\sin z^2$ or $\sin^2 z$ would yield for $\sigma_i(x, y)$ the function $\sin x \cosh y$, $\sin(x^2 - y^2) \cosh(2xy)$, $\sin^2 x \cosh^2 y - \cos^2 x \sinh^2 y$ and $\sin(\sin x \cosh y) \cosh(\cos x \sinh y)$, respectively.

Combining these choices of basis functions with linear transformations as defined in Eq. equation 6
 with Eq. equation 5 yields a single hidden layer superposition network constrained to Laplace's equation.

240 **Remark 1** (The universality of superposition networks) As the weighted summation of the real 241 parts of holomorphic functions, it follows that the Laplacian superposition network as outlined is 242 the real part of a holomorphic function. Such functions cannot represent arbitrary harmonic functions. In the case of Ω being a multiply-connected domain, there are harmonic functions that cannot 243 be represented as the real-part of a holomorphic function. There are means to circumvent such re-244 strictions (Ghosh et al., 2023). However, the wider topic of proving the universality of a function 245 class within a space of solutions of arbitrary linear differential operators is, to our knowledge, yet 246 an open area of study. We thus leave investigations of universal superposition networks to future 247 work. 248

249 2.2.2 DIVERGENCE-FREE FIELDS

Similarly to the Laplace equation constrained architectures of section 2.2.1, divergence-free fields, given in the scenario where $\mathcal{L}\phi = \nabla \cdot \phi$, obey the same symmetries as defined in Eq. equation 6.

It remains to define suitable activation functions. Previous work on deriving divergence-free neural network architectures (Ghosh et al., 2023; Richter-Powell et al., 2022) has demonstrated this approach.

In two-dimensions, given a differentiable function $f : \mathbb{R}^2 \to \mathbb{R}$, we can derive a divergence-free field $g : \mathbb{R}^2 \to \mathbb{R}^2$ via the following transformation:

$$(\hat{x}, \hat{y}) = \left(\frac{\partial f(x, y)}{\partial y}, -\frac{\partial f(x, y)}{\partial x}\right).$$
(7)

The explicit forms used to derive superposition networks in our numerical experiments are enumerated in section 3.4. In three-dimensions, given a differentiable function $f : \mathbb{R}^3 \to \mathbb{R}^3$, we can derive a divergence-free field $g : \mathbb{R}^3 \to \mathbb{R}^3$ using:

$$(\hat{x}, \hat{y}, \hat{z}) = \nabla \times f(x, y, z) \tag{8}$$

2.2.3 HEAT EQUATION

We consider architectures constrained to satisfy the 2+1 (two spatial and one temporal) dimensional heat equation, where $\mathcal{L} = \frac{\partial}{\partial t} - \alpha \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right)$ in Eq. equation 1, where α is the thermal diffusivity coefficient. 270 In contrast to section 2.2.1, which constructed generally applicable Lie group symmetries of the 271 Laplace equation, here we demonstrate the construction of a superposition network for an appropri-272 ate Lie group via a manufactured solution of the heat equation.

273 Consider functions $\sigma_i(x, y, t)$ of the following form: 274

$$\sigma_i(x, y, t) = e^{-2\alpha t} \sin x \sin y, \tag{9}$$

276 which provides a manufactured solution to the heat equation suitable as the activation function in 277 Eq. equation 5.

278 Given the heat-equation and the chosen activation function, it can be validated through direct sub-279 stitution that the following group action provides a trainable transformation $(x, y, t) \rightarrow (\hat{x}, \hat{y}, \hat{t})$ for 280 use as the linear component of Eq. equation 5: 281

282 283

284 285

298 299 300

301

275

 $\hat{x_i} = \theta_{x1i}x + \theta_{x2i}$ $\hat{y}_i = \theta_{y1i}y + \theta_{y2i}$ (10) $\hat{t_i} = \frac{\theta_{x1i}^2 + \theta_{y1i}^2}{2}t + \theta_{ti},$

286 whence it follows ¹ that $\mathcal{L}\sigma_i(\hat{x}_i, \hat{y}_i, \hat{t}_i) = 0$, and hence constraining the linear layer of a superposi-287 tion network to follow Eq. equation 10, with an activation function given by Eq. equation 9, yields 288 an architecture automatically satisfying the heat equation. 289

290 2.2.4 **BURGERS' EQUATION**

291 So far, we investigated how the construction of superposition networks can constrain architectures to 292 certain linear differential operators, leaving unaddressed the case of nonlinear differential equations. 293 However, the latter can be converted to (and from) linear differential equations through appropriate transformations in certain important cases. Applying such transformations to superposition networks 295 thus extends the utility of these architectures to cases of nonlinear differential equations. 296

To this end, we consider 1D Burgers' equation: 297

$$\frac{u}{\partial t} + u\frac{\partial u}{\partial x} = \nu \frac{\partial^2 u}{\partial x^2} \tag{11}$$

If $\nu \neq 0$, we can apply the Cole-Hopf transformation (Cole, 1951; Hopf, 1950): u(x,t) = $-2\nu \frac{\partial}{\partial x} \ln \phi(x,t)$ which yields the heat equation

$$\frac{\partial \phi}{\partial t} = \nu \frac{\partial^2 \phi}{\partial x^2} \tag{12}$$

If $\phi(x,t)$ is defined via the means outlined in section 2.2.3 with Eq. equation 5, it follows that 306 $\phi(x,t)$ naturally satisfies the heat equation. We can thus train the superposition network on the 307 boundary and initial conditions inherited from the Burgers' equation and finally apply the inverse 308 transformation, where the derivatives can be calculated readily via automatic differentiation. The resulting architecture is then guaranteed to be constrained to Burgers' equation.

310 This approach can be generalized to 2D and 3D Burgers' equations 311

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = \nu \nabla^2 \mathbf{u}$$
(13)

313 via the transformation $\mathbf{u} = -2\nu\nabla \ln \phi$. 314

315 316

317

312

309

3.1 Methods 318

319 There are two principal manners in which to provide a differential-equation based inductive bias 320 to a neural network: 1. with appropriate regularisation terms, such as physics-informed neural net-321 works(Raissi et al., 2019). 2. by incorporating differential equation constraints within the architec-322 ture itself (Ghosh et al., 2023; Richter-Powell et al., 2022). 323

¹Note how the partial derivatives of \mathcal{L} are still taken with respect to the original coordinates of (x, y, t).

We seek to benchmark superposition networks against both approaches where possible, but limit ourselves to PINNs where no such hard-constrained architecture exists.

General Formulation We consider differential equations as outlined in equation 1, however, we allow \mathcal{L} to also be nonlinear to encompass our handling of Burgers' equation.

Introduce probability distributions of collocation points over the interior of a domain \mathbb{P}_{Ω} and the boundary of the domain $\mathbb{P}_{\partial\Omega}$.

PINN(BI) PINNs use appropriately constructed regularisation terms to provide inductive biases towards a given differential equation. Many neural network architectures are amenable to physics-informed regularisation, yet MLPs tend to be the most widely used. To optimise a PINN to solve a forward-solution of a given differential, the following loss function is minimized:

$$L_{\text{PINN}}(\theta) = \mathbb{E}_{x \sim \mathbb{P}_{\partial \Omega}} \left[\left(\mathcal{N} f_{\theta}(x) \right)^2 \right] + \mathbb{E}_{x \sim \mathbb{P}_{\Omega}} \left[\left(\mathcal{L} f_{\theta}(x) \right)^2 \right], \tag{14}$$

where f_{θ} represents an MLP with trainable parameters and the two terms represents contributions of the boundary and interior respectively. We denote by PINNB/PINNI methods which provide an extra weight (by a factor of 1000 in our numerical experiments) to the boundary and interior loss terms, respectively.

RAR The expectation over Ω in equation 14 is often done over a fixed distribution. In Lu et al. (2021a), an adaptive approach is introduced whereby points with high PDE residuals are adaptively added to the pool of points to estimate of the second term of equation 14. We name such approach RAR and include it within our benchmarks: every 1000 epochs of training, we sample 1024 new candidate collocation points, and add the 32 with the highest PDE residuals to the pool of collocation points used for each subsequent training epochs.

AA Adaptive activation functions were proposed by Jagtap et al. (2020) to improve convergence of PINNs. In particular, they propose that given an MLP $f(x) = W_L(\sigma(W_{L-1}(...(\sigma(W_1x)))))$, that it be modified to be of the form $f(x) = W_L(an\sigma(W_{L-1}(...(an\sigma(W_1x)))))$ for a trainable parameter *a* and fixed hyperparameter *n*. Following the original publication, we initialise *a* = 1.0 and fix *n* = 10 for all our experiments. We also include a combination of this technique with the residual adaptive approach which we denote as **RAR+AA**.

Holomorphic Neural Networks We implement holomorphic neural networks as a means to model the Laplace equation, following Ghosh et al. (2023).

Holomorphic neural networks parameterise a holomorphic function with a complex-valued MLP $f_{\theta}(x + iy) : \mathbb{C} \to \mathbb{C}$. If the activation functions of the MLP are themselves holomorphic, then the real part of the output is guaranteed to satisfy Laplace's equation in two-dimensions with inputs xand y. A holomorphic neural network can thus be trained to solve instances of Laplace's equations in two-dimensions by minimising the following loss function.

$$L_{\text{Holomorphic}}(\theta) = \mathbb{E}_{x \sim \mathbb{P}_{\partial \Omega}} \left[\mathcal{N} \left(\text{Re}(f_{\theta}(x)) \right) \right], \tag{15}$$

where $\operatorname{Re}(z)$ denotes taking the real part of a complex number z and f_{θ} represents a holomorphic neural network. In all numerical experiments, we use sin as our holomorphic activation function.

NCL It is possible to derive divergence-free neural networks in arbitrary dimensions (Richter Powell et al., 2022). However, in our numerical experiments, we limit ourselves to the two dimensional case, which allows us to simplify our exposition compared to the more general for mulations presented in (Richter-Powell et al., 2022).

Define a neural network $f_{\theta} : \mathbb{R}^2 \to \mathbb{R}$ as a multilayer perceptron with trainable parameters θ . Then we can postprocess f_{θ} as follows to achieve a divergence-free function $g_{\theta} : \mathbb{R}^2 \to \mathbb{R}^2$ given by applying Eq. equation 7 to f_{θ} . We optimise divergence-free neural networks as per the loss function in Eq. equation 14.

375

363

366

327

328

337

348

Experimental Details Details common for every numerical experiment are outlined in this subsection. All experiments reported in Table 1 were executed in Python3, making use of NumPy (Harris et al., 2020) and Matplotlib (Hunter, 2007) libraries, with all the neural networks implemented

Table 1: A summary of experimental results (lower is better) comparing superposition networks 379 to alternative architectures imposing differential equation constraints for the methods outlined in 380 section 3.1. Root-mean squared errors (RMSEs) of the final trained solution are shown with standard 381 deviations over 10 random seeds reported. For the heat equation, we report the RMSE at the end of 382 the simulation. Note that Holomorphic neural networks are only applicable to Laplace's equation, 383 and NCL only applies to divergence-free fields. 384

95		Laplace 1	Laplace 2	Heat 1	Heat 2	Navier Stokes	Burgers'
05	Superposition	0.0067 ± 0.0023	$0.010 {\pm} 0.0047$	0.0080±7.3e-5	$0.00084{\pm}0.00018$	0.11±0.0026	0.0030±0.0024
86	PINN	$0.15 {\pm} 0.0030$	$0.29 {\pm} 0.093$	$0.0085 {\pm} 0.0024$	$0.0027 {\pm} 0.0017$	$0.10{\pm}0.00090$	$0.0039 {\pm} 0.0022$
87	PINNB	$0.0063 {\pm} 0.0041$	$0.12{\pm}0.066$	$0.063 {\pm} 0.016$	$0.015 {\pm} 0.0025$	$0.085 {\pm} 0.0090$	$0.0049 {\pm} 0.0030$
000	PINNI	$0.56 {\pm} 0.00091$	$0.82{\pm}0.067$	$0.080{\pm}0.024$	0.046±4.4e-5	$0.10{\pm}0.00034$	$0.12{\pm}0.010$
88	RAR	$0.15 {\pm} 0.0015$	0.43 ± 0.012	$0.0085 {\pm} 0.0075$	0.0026 ± 0.0012	$0.10{\pm}0.00058$	$0.0036 {\pm} 0.00065$
89	AA	$0.19{\pm}0.12$	$0.54{\pm}0.084$	$0.039 {\pm} 0.015$	0.016 ± 0.016	$0.097{\pm}0.00086$	0.0067 ± 0.0039
0.0	RAR+AA	$0.20{\pm}0.12$	$0.55 {\pm} 0.063$	$0.0070 {\pm} 0.0020$	$0.0053 {\pm} 0.0034$	$0.099 {\pm} 0.00080$	$0.018 {\pm} 0.012$
90	NCL	-	-	-	-	$0.097 {\pm} 0.0015$	-
91	Holomorphic	0.0029 ± 0.0022	$0.0033 {\pm} 0.0009$	-	-	-	-

392 393

416 417 418

421

378

via PyTorch (Paszke et al., 2019). In the supplementary material we also report an independent 394 implementation of the PINN baseline for the Laplace benchmark using JAX Bradbury et al. (2018), 395 which does not materially differ from the numerical results presented. 396

397 Ground truths for the heat equation and Navier-Stokes benchmarks were constructed using FEATools commercial software with the open-source OpenFOAM backend (Weller et al., 1998) 398 for Navier-Stokes, and MATLAB's backslash backend for the heat equation (Amestoy et al., 2000). 399 While the source-code is based on proprietary software, we include CSVs of simulation output for 400 use with our analyses for the purpose of reproducibility along with the source code of all experi-401 ments in the supplementary material. We consistently use Kaiming Uniform initialisers (He et al., 402 2015), optimised for 32000 epochs over full-batches with an Adam optimizer (Kingma & Ba, 2014) 403 with a learning rate (LR) of 10^{-3} . We use tanh activations for real-valued neural networks and 404 sin activations for holomorphic networks. All experiments run using Python 3.10 on two-cores of a 405 Dual AMD Rome 7742 processor with 8GB of RAM and were allocated 12 hours of compute time, 406 but finished well-within that period.

407 In the following sections, we present the specific methods and results used for each differential 408 equation. All experiments are repeated for 10 different random seeds with means and standard 409 deviations of all results shown in Table 1. 410

411 3.2 LAPLACE'S EQUATION 412

413 Consider a domain given by $\Omega = (0,1) \times (0,1)$, with $\partial \Omega$ corresponding to the boundary lines at 414 x = 0, y = 0, x = 1, y = 1.415

We construct a target function given by:

$$f(x,y) = \operatorname{Re}\left\{\frac{1}{(z-1.2-0.5i)(z+0.2-0.5i)(z-0.5+0.2i)(z-0.5-1.2i)}\right\},$$
 (16)

419 where z = x + iy. Note that we construct a holomorphic function in Eq. equation 16 that is not 420 holomorphic throughout the entirety of \mathbb{C} . This prevents the solution of the function from appearing as a trivial solution of the superposition and holomorphic neural networks. 422

As boundary conditions, we take Dirichlet conditions of f(x, y) on $\partial \Omega$ for the benchmark that we 423 refer to as Laplace 1, and we take Neumann boundary conditions in the y-direction at y = 1 for the 424 case of Laplace 2. 425

426 For the superposition networks, we take as activation functions the real parts of the following func-427 tions $\sin z$, $\sin z^2$, $\sin^2 z$, $\sin \sin z$, e^z , $e^{\sin z}$, $\sin e^z$ and e^{z^2} as our chosen activation functions in Eq. equation 5, with 32 repeats of functions each. We implement elements of the Euclidean sym-428 metry group in Eq. equation 6 as a rotation followed by a translation, with half of them mirrored 429 to allow for orientation to be preserved or flipped. Initial angles of rotations are sampled uniformly 430 on $(0, 2\pi)$. Shift amounts in Eq. equation 6 are initialised to zero. All other parameters in the su-431 perposition network initialised to zero. For holomorphic neural networks and PINNs, we use MLPs with three hidden layers of width 64. We provide training points on 128 evenly spaced points on each of the lines x = 0, x = 1, y = 0, y = 1 for collocation points on the boundary, and on 1024 uniformly-random distributed points on the interior of Ω .

We find that neural network architectures constrained to follow the Laplace equation tend to perform better than PINNs in our benchmarks, as demonstrated in table 1. Holomorphic neural networks maintain a strong performance across both benchmarks, surpassing superposition networks. However, holomorphic neural networks are only applicable towards modelling Laplace's equation in 2D,wwhereas the superposition network methodology applies to arbitrary linear differential equations in arbitrary dimensions.

3.3 HEAT EQUATION

444 We consider two heat equation benchmarks in two-spatial dimensions with different initial condi-445 tions:

$$\phi(x,y,0) = \sqrt{e^{-5((x-0.5)^2 + (y-0.5)^2)}(\sin^2 5\pi x + \cos^2 3\pi y)}$$
(17)

and

435

442

443

446 447

448 449 450

451

453 454

455

$$\phi(x, y, 0) = e^{-5((x-0.5)^2 + 10(y-0.5)^2)} - e^{-20((x-0.5)^2 + 5(y-0.7)^2)} - e^{-20((x-0.5)^2 + 5(y-0.3)^2)}$$
(18)

for what we refer to as Heat1 and Heat2 in table 1, respectively.

452 As for boundary conditions, we take mixed Dirichlet and Neumann boundary conditions such as

$$\phi(0, y, t) = \phi(1, y, t) = 0 \qquad \left. \frac{\partial \phi}{\partial y} \right|_{y=0} = \left. \frac{\partial \phi}{\partial y} \right|_{y=1} = 0 \tag{19}$$

for Heat1 (17) and

$$\frac{\partial \phi}{\partial x}\Big|_{x=0} = \left.\frac{\partial \phi}{\partial x}\right|_{x=1} = 0 \qquad \phi(x,0,t) = \phi(x,1,t) = 0 \tag{20}$$

460 for Heat2 (18).

As with the Laplace equation in section 3.2, we are careful to ensure that the initial conditions are not directly representable by a single term of a superposition, thus many superposition bases are required to approximate the resulting solution.

In this instance, we only benchmark against PINNs since NCL and holomorphic neural networks do not represent the heat equation.

For the superposition networks, we implement architectures as per section 2.2.3 with 64 trainable components. We initialise all parameters in Eq. equation 10 and Eq. equation 5 random-uniformly on (0, 1), with the exception of θ_{x1i} and θ_{y1i} , which are initialised random-uniformly on (0, 10). We use MLPs with 3 hidden layers of width 64 for PINNs.

We find that the superposition network methodology is able to perform better than the PINN on
both problems, even though the approximation of the initial solution is not as precise. The diffusion
phenomenon is, however, very well captured, proving that the proposed method can handle the extra
layer of complexity brought by the time dependency of the differential equation.

475 476

480

481

3.4 INCOMPRESSIBLE NAVIER-STOKES EQUATION

In this section, we demonstrate the challenges remaining in physics-informed optimisation for practical problems. We consider solutions to the incompressible steady-state Navier-Stokes equation:

$$(\mathbf{u} \cdot \nabla) \,\mathbf{u} = \nu \nabla^2 \mathbf{u} - \frac{1}{\rho} \nabla \mathbf{p},\tag{21}$$

with the additional constraint that $\nabla \cdot u = 0$ everywhere. We denote the x and y components of **u** with unbolded symbols u and v, respectively.

We take Ω as $(0,0.5) \times (0,0.1)$, with two circles of radius 0.05 removed from the rectangle. The centres of the circles lie at (0.15,0.1) and (0.35,0.0). As boundary conditions, we take $\mathbf{u} = (0.1,0)$

486 487 488 at x = 0, $\mathbf{u} = w$ on the circles, p = 0 at x = 0.5, and v = 0 elsewhere. The constraint of v = 0corresponds to imposing a reflection of the velocity field parallel to the x-axis at y = 0.15 and y = 0, thus this represents a tiled structure.

489 For the superposition network setups, we follow to use the following forms for σ_i in Eq. equation 5: 490 $(\cos(x+y), -\cos(x+y)), (e^{x+y}, -e^{x+y}), (x\cos(xy), -y\cos(xy)))$. Our initialisation and num-491 ber of symmetry groups follow the same parameters as per section 3.2. For PINNs, NCL and the 492 networks representing p we use MLPs with three hidden layers of width 32 each. The PINN has an 493 additional loss term for divergence-free constraints. We sample equally spaced points on $\partial \Omega$ such 494 that each line segment has a density of 1000 collocation points per unit length, with the exception of 495 each semi-circle which has 100 collocation points on the semi-circle each. The interior collocation 496 points are formed with rejection sampling with 1024 random uniformly distributed points distributed over the rectangle $(0, 0.5) \times (0, 0.1)$. 497

We find it noteworthy that in this scenario where nonlinear Navier-Stokes constraints cannot be embedded within the architecture, convergence is problematic for all methods.

501 3.5 BURGERS' EQUATION

503 We construct a benchmark defined as per the 1D-Burgers' equation in equation 11 with $\nu = 0.1$, 504 initial conditions of $u(x, 0) = \exp(-50(x-0.6)^2) - \exp(-50(x-0.4)^2)$ and boundary conditions 505 of u(0, t) = 0 and u(1, t) = 1, with $\Omega = (0, 1)^2$.

For superposition networks, we consider 100 components with parameters chosen analogously to
 the heat equation experiments of section 3.3. We use 64 evenly spaced points on each Dirichlet
 boundary for imposing boundary conditions and 1024 collocation points on the interior for PINN
 based methods.

510 511

4 **DISCUSSION**

512 513

Motivated by the difficulty of training physics-informed regularisation, this work has sought to em-514 bed linear differential equation constraints within neural network architectures, and also theoreti-515 cally demonstrated the possibility of extending this to some non-linear differential equations. We 516 suggested a systematic approach to further extend our approach to include additional differential 517 operators, and hence ODE/PDEs, as well as illustrating ad-hoc constructions for notable cases. Cru-518 cially, ensuring convergence guarantees for equations at the interior of the domain might enable 519 the adoption of neural architectures in more critical applications than is currently targeted. Numer-520 ical investigations support a perspective that favours embedding differential equations constraints 521 via architectural design rather than regularisation. However, we emphasise some limitations in our 522 approaches. While we present Lie-group symmetries and manufactured solutions of certain linear-523 differential equations, we do not present methodical ways to derive such solutions for arbitrary linear differential equations. In practice, for other differential equations, a practitioner might em-524 ploy techniques such as a trial ansatz, using physically-motivated arguments, or attempting standard 525 approaches used in nonlinear differential equations (Hydon, 2000). We hope that our work moti-526 vates further research into embedding differential equation constraints directly into neural network 527 architectures. 528

529 530

531

532

533 534

535

536

References

Patrick R Amestoy, Iain S Duff, Jean-Yves L'Excellent, and Jacko Koster. Mumps: a general purpose distributed memory sparse solver, 2000.

Gert-Jan Both, Subham Choudhury, Pierre Sens, and Remy Kusters. Deepmod: Deep learning for model discovery in noisy data. *Journal of Computational Physics*, 428:109985, 2021.

 James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal
 Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao
 Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL http: //github.com/jax-ml/jax. 551

567

569

570

571

583

- 540 Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary 541 differential equations. 31:6571–6583, 2018. 542
- Julian D. Cole. On a quasi-linear parabolic equation occurring in aerodynamics. *Quarterly of* 543 Applied Mathematics, 9:225-236, 1951. URL https://api.semanticscholar.org/ 544 CorpusID: 39662248.
- 546 Atiyo Ghosh, Antonio Andrea Gentile, Mario Dagrada, Chul Lee, Seong-Hyok Sean Kim, 547 Hyukgeun Cha, Yunjun Choi, Dongho Kim, Jeong-Il Kye, and Vincent Emanuel Elfving. Harmonic neural networks. pp. 11340-11359, 2023. 548
- 549 Robert J Gray. How to calculate all point symmetries of linear and linearizable differential equa-550 tions. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 471(2175):20140685, 2015. 552
- Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. 32:15379-553 15389, 2019. 554
- 555 Marc Harkonen, Markus Lange-Hegermann, and Bogdan Raita. Gaussian process priors for systems 556 of linear partial differential equations with constant coefficients. pp. 12587–12615, 2023.
- Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David 558 Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti 559 Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, 561 Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. Nature, 585:357-362, 2020. 563
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing 564 human-level performance on imagenet classification. pp. 1026–1034, Dec 2015. 565
- 566 Eberhard Hopf. The partial differential equation $u_t + uu_x = \mu u_{xx}$. Communications on Pure and Applied Mathematics, 3(3):201–230, 1950. doi: https://doi.org/10.1002/cpa.3160030302. URL 568 https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.3160030302.
 - J. D. Hunter. Matplotlib: A 2d graphics environment. Computing in Science & Engineering, 9(3): 90-95, 2007.
- 572 Peter Ellsworth Hydon. Symmetry methods for differential equations: a beginner's guide. Num-573 ber 22. Cambridge University Press, 2000.
- 574 Ameya D Jagtap and George Em Karniadakis. Extended physics-informed neural networks (xpinns): 575 A generalized space-time domain decomposition based deep learning framework for nonlinear 576 partial differential equations. Communications in Computational Physics, 28(5), 2020. 577
- Ameya D Jagtap, Kenji Kawaguchi, and George Em Karniadakis. Adaptive activation functions 578 accelerate convergence in deep and physics-informed neural networks. Journal of Computational 579 Physics, 404:109136, 2020. 580
- 581 Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint 582 arXiv:1412.6980, 2014.
- Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W Mahoney. Char-584 acterizing possible failure modes in physics-informed neural networks. 34:26548–26560, 2021. 585
- 586 Isaac E Lagaris, Aristidis Likas, and Dimitrios I Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. IEEE Transactions on Neural Networks, 9(5):987-1000, 1998. 588
- 589 Lu Lu, Xuhui Meng, Zhiping Mao, and George Em Karniadakis. Deepxde: A deep learning library 590 for solving differential equations. SIAM review, 63(1):208-228, 2021a. 591
- Lu Lu, Raphael Pestourie, Wenjie Yao, Zhicheng Wang, Francesc Verdugo, and Steven G Johnson. 592 Physics-informed neural networks with hard constraints for inverse design. SIAM Journal on Scientific Computing, 43(6):B1105–B1132, 2021b.

594 595	Zhiping Mao, Ameya D Jagtap, and George Em Karniadakis. Physics-informed neural networks for high-speed flows. <i>Computer Methods in Applied Mechanics and Engineering</i> , 360:112789, 2020.			
596 597	Ben Moseley, Andrew Markham, and Tarje Nissen-Meyer. Finite basis physics-informed neural networks (fbpinns): a scalable domain decomposition approach for solving differential equations. <i>Advances in Computational Mathematics</i> , 49(4):62, 2023.			
590 599				
600	Saviz Mowlavi and Saleh Nabi Optimal control of pdes using physics-informed neural networks			
601	Journal of Computational Physics, 473:111731, 2023.			
602	Mohammad Amin Nabian and Hadi Meidani. A deep neural network surrogate for high-dimension random partial differential equations. <i>arXiv preprint arXiv:1806.02957</i> , 2018.			
603 604				
605 606	uofei Pang, Lu Lu, and George Em Karniadakis. fpinns: Fractional physics-informed neural networks. <i>SIAM Journal on Scientific Computing</i> , 41(4):A2603–A2626, 2019.			
607 608 609 610 611 612	Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. 32:8024–8035, 2019.			
613 614 615	Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. <i>Journal of Computational physics</i> , 378:686–707, 2019.			
616 617	ack Richter-Powell, Yaron Lipman, and Ricky TQ Chen. Neural conservation laws: A divergence- free perspective, 35:38075–38088, 2022.			
618 619 620	Hwijae Son, Sung Woong Cho, and Hyung Ju Hwang. Enhanced physics-informed neural networks with augmented lagrangian relaxation method (al-pinns). <i>Neurocomputing</i> , 548:126424, 2023.			
621 622 623	Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. <i>arXiv preprint arXiv:2011.13456</i> , 2020.			
625 626 627	Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. <i>SIAM Journal on Scientific Computing</i> , 43(5):A3055–A3081, 2021.			
628 629 630 631	enry G Weller, Gavin Tabor, Hrvoje Jasak, and Christer Fureby. A tensorial approach to com- putational continuum mechanics using object-oriented techniques. <i>Computers in physics</i> , 12(6) 620–631, 1998.			
632				
633				
634				
635				
636				
637				
630				
640				
641				
642				
643				
644				
645				
646				
6/17				