


GLENN: Neural network-enhanced computation of Ginzburg–Landau energy minimizers

Michael Crocoll, Christian Döding,
Benjamin Dörich, Roland Maier

CRC Preprint 2026/11, April 2026

KARLSRUHE INSTITUTE OF TECHNOLOGY

CRC 1173



Participating universities



Funded by



GLENN: NEURAL NETWORK-ENHANCED COMPUTATION OF GINZBURG–LANDAU ENERGY MINIMIZERS

MICHAEL CROCOLL*, CHRISTIAN DÖDING†, BENJAMIN DÖRICH*, AND ROLAND MAIER*

Abstract. In this work, we propose a neural network-enhanced finite element strategy to compute the minimizer of the Ginzburg–Landau energy based on an unsupervised deep Ritz-type strategy. We treat the parameter κ as a variable input parameter to obtain possible minimizers for a large range of κ -values. This allows for two possible strategies: 1) The neural network may be extensively trained to work as a stand-alone solver. 2) Neural network results are used as starting values for a subsequent classical iterative minimization procedure. The latter strategy particularly circumvents the missing reliability of the neural network-based approach. Numerical examples are presented that show the potential of the proposed strategy.

Key words. Ginzburg–Landau equation, superconductivity, energy minimization, finite element method, neural networks.

MSC codes. 35Q56, 65K10, 68T07, 65N30

1. Introduction. In 1911, the Dutch physicist Onnes [23] realized that when mercury is cooled below some critical temperature, the electrical resistance vanishes, i.e., the mercury is in a superconducting state. At the same time, one can observe that an applied magnetic field below some critical strength is expelled from the superconducting material. On the other hand, above a critical strength, the magnetic field penetrates and destroys the superconducting property. However, for so-called type-II superconductors, there is an intermediate regime in which the magnetic field only penetrates at isolated points while leaving the main part of the material in the superconducting state. Such a property is of particular interest in applications where the superconductor then becomes robust to weak perturbative magnetic fields.

The mathematical model to describe the phenomena of superconductivity on a macroscopic level is the Ginzburg–Landau (GL) model [13, Sec. 3]. Let $\Omega \subset \mathbb{R}^d$ with $d = 2$ or $d = 3$ denote the region occupied by the superconducting material. The superconducting properties are described by a complex-valued function $u : \Omega \rightarrow \mathbb{C}$, called the *order parameter*. Although u itself is not directly observable, similar to a quantum mechanical wave function, its squared modulus $|u|^2$ can be measured in physical experiments and interpreted as the density of superconducting electrons (*Cooper pairs*). The value of $|u|^2$ ranges between 0 and 1 where $|u(x)|^2 = 1$ belongs to regions of superconductivity and $|u(x)|^2 = 0$ to regions where the material is in its normal (non-superconducting) state. Superconductivity is rare in nature but has a wide range of possible applications.

Beyond the fully superconducting state ($|u|^2 \equiv 1$) and the normal state ($|u|^2 \equiv 0$), mixed states with $0 \leq |u(x)|^2 \leq 1$ arise, in particular, when an external magnetic field \mathbf{h}_{ext} is applied. The mixed normal-superconducting state can include a co-existence of superconducting and normal states – a so-called *Abrikosov vortex lattice* [1], where the density is zero in vortex centers. Such configurations are considered in this work, which, however, can only occur for special superconductors with appropriate magnetic

*Institute for Applied and Numerical Mathematics, Karlsruhe Institute of Technology, 76149 Karlsruhe, Germany (michael.crocoll@kit.edu, benjamin.doerich@kit.edu, roland.maier@kit.edu).

†Institute for Numerical Simulation, University of Bonn, 53115 Bonn, Germany (doeding@ins.uni-bonn.de).

fields. In particular, the internal magnetic field, represented by $\operatorname{curl} \mathbf{A}$, tends to align with \mathbf{h}_{ext} , where $\mathbf{A} : \Omega \rightarrow \mathbb{R}^d$ is the (typically unknown) magnetic vector potential within the superconductor. From a mathematical perspective, for the state of the superconductor the GL model prescribes that (u, \mathbf{A}) minimizes the *GL free energy* (cf. [13, Sec. 3]) given by

$$(1.1) \quad E(u, \mathbf{A}) := \frac{1}{2} \int_{\Omega} \left| \frac{i}{\kappa} \nabla u + \mathbf{A}u \right|^2 + \frac{1}{2} (1 - |u|^2)^2 + |\operatorname{curl} \mathbf{A} - \mathbf{h}_{\text{ext}}|^2 dx,$$

where $\kappa > 0$ is a material parameter, also called the *GL parameter*.

As indicated above, one is interested in vortex states, where κ is linked to the structure of the vortex lattice, see [27, 29, 30, 31]. There are no vortices for small values of κ , while larger values lead to more and more (very localized) vortices [3, 27]. High values of κ are particularly interesting and challenging, considering especially the numerical approximation of corresponding minimizers.

The numerical approximation of minimizers of (1.1) has first been studied in [13, 14]. Therein, error estimates in H^1 have been derived for both the order parameter u and the magnetic vector potential \mathbf{A} when using a finite element discretization. The convergence results, however, do not take into account the role of κ , which can generally lead to constraints on the mesh size as investigated in [12] for a reduced setting with a fixed potential \mathbf{A} . In that setting, the energy in (1.1) is only minimized with respect to u . The derived constraints on the choice of the mesh size may be overcome by the use of appropriate multiscale constructions as realized in [6]. The full problem is considered in [9], where different meshes and spaces for u and \mathbf{A} are used as an adaptation to the different scaling behavior in κ . More recently, in [7] the precise coupling of not only the spatial mesh width h and the parameter κ , but also the polynomial degree p is studied. The authors are able to precisely quantify the pre-asymptotic regime and the corresponding pollution effect, and further show how higher-order polynomials reduce these effects. An extension of this technique to multiscale approximation spaces is presented in [10]. A generalization to non-convex domains is provided in [11]. Here, Nédélec elements and a discrete divergence-free condition have to be used to anticipate the low regularity of the vector potential. This is also the framework used in this work. Besides, we mention here the works [16] and [15], which conduct an error analysis for a covolume and a finite volume method, respectively (without explicit rates in the mesh size or κ).

All of these works focus on the approximability of minimizers. For the practical computation in the numerical examples, different choices of minimization procedures are employed: in [6] and [12], a semi-implicit L^2 -gradient flow is used, whereas in [9] a Sobolev and in [7, 10] a conjugate Sobolev gradient flow approach are employed. Even though it becomes clear that it appears to be always preferable to apply a conjugate Sobolev gradient flow for the GL problem, the results are not satisfying in the sense of reliably capturing the best minimizers without trial and error to find good starting values. An alternative approach is to use physics-based neural networks to approximate minimizers. In particular, the GL energy could be naturally used as a loss function in the spirit of a deep Ritz method [17] that is to be minimized by the training of a neural network. Nonetheless, a pure machine learning-based strategy can also be very sensitive to small changes in the network architecture.

In this work, we propose a neural network-enhanced strategy for the approximation of minimizers (of the discrete version) of (1.1). An essential issue of classical iterative algorithms is the choice of a starting value, which can heavily influence the number of required iterations and even the actual limit. To overcome this, we pro-

pose to train a neural network based on a certain parameter range of κ to obtain a reasonable initial guess that leads to better GL minimizers. In essence, this also circumvents the problematic properties of neural network-based schemes, which to date are not fully reliable in the sense of a provable convergence behavior. More precisely, even if the suggested initial guess is not optimal, the subsequent classical iterative process still follows deterministic and theoretically grounded principles. Nonetheless, we also consider the setting where a network is trained to directly output a solution pair (u, \mathbf{A}) for a chosen value of κ . This, however, requires a certain amount of pre-computations to set up a valid model. The latter approach is particularly meant to show that alternatives to classical iterative algorithms based on finite element spaces have the potential to lead to even better minimizers.

There exist similar approaches in the literature, where classical algorithms are enhanced by neural networks. Recent examples are in the context of nonlinear Schrödinger eigenvalue problems [24] or to speed up iterative algorithms for convex optimization [34] or inverse problems [2]. However, to the best of our knowledge they all use supervised learning (i.e., training with data), while our approach is unsupervised, i.e., it does not need data and solely relies on the energy-based loss function for the training.

The rest of the paper is structured as follows. In Section 2, we recall the most important results on the minimizer of the GL energy and explain the established methods and how to derive our neural network-enhanced method from it. We present the neural networks approach in Section 3 and explain the finite element solver in Section 4. Finally, we provide numerical experiments in Section 5.

2. Ginzburg-Landau problem and discrete minimizers.

2.1. Preliminaries. For the mathematical framework, we require the following notation: For $2 \leq p \leq \infty$ we denote by $L^p = L^p(\Omega, \mathbb{C})$ the standard Lebesgue space of scalar complex-valued function on Ω equipped with the usual norm $\|\cdot\|_{L^p}$. For $p = 2$ the Hilbert space L^2 is considered as a *real* Hilbert space with the inner product $(v, w) = \operatorname{Re} \int_{\Omega} v \bar{w} dx$, where \bar{w} denotes the complex conjugate of w . Similarly, we denote by $\mathbf{L}^p = L^p(\Omega, \mathbb{R}^d)$ the space of p -Lebesgue-integrable vector fields, with norm $\|\cdot\|_{\mathbf{L}^p}$, and the inner product on \mathbf{L}^2 is given by $(\mathbf{B}, \mathbf{C}) = \int_{\Omega} \mathbf{B} \cdot \mathbf{C} dx$. Note, that we use the same notation (\cdot, \cdot) for both inner products on L^2 and \mathbf{L}^2 as its meaning is clear from the context. Accordingly, we denote by $H^1 = H^1(\Omega, \mathbb{C})$ with norm $\|\cdot\|_{H^1}$ the Sobolev space of complex-valued functions whose weak derivative exists in L^2 . The same notation applies to vector fields and we write $\mathbf{H}^1 = H^1(\Omega, \mathbb{R}^d)$. For $d = 3$, we denote by $\mathbf{H}(\operatorname{curl})$ the space of vector fields with weak rotation in \mathbf{L}^2 and by $\mathbf{H}(\operatorname{div})$ the space of vector fields with weak divergence in L^2 . Both spaces can be equipped with natural zero boundary conditions leading to the subspaces

$$\begin{aligned} \mathbf{H}_0(\operatorname{curl}) &= \{\mathbf{B} \in \mathbf{H}(\operatorname{curl}) \mid \mathbf{B} \times \nu = 0 \text{ on } \partial\Omega\}, \\ \mathbf{H}_0(\operatorname{div}) &= \{\mathbf{B} \in \mathbf{H}(\operatorname{div}) \mid \mathbf{B} \cdot \nu = 0 \text{ on } \partial\Omega\}, \end{aligned}$$

where ν denotes the outer unit normal on $\partial\Omega$. Finally, we set

$$\mathbf{H}(\operatorname{curl}, \operatorname{div}) = \mathbf{H}(\operatorname{curl}) \cap \mathbf{H}(\operatorname{div}), \quad \mathbf{H}_0(\operatorname{curl}, \operatorname{div}) = \mathbf{H}(\operatorname{curl}) \cap \mathbf{H}_0(\operatorname{div}).$$

For $d = 2$, there are two curl operators that act either on scalar fields or on vector fields. For a scalar field $\mathbf{h}_{\text{ext}} = h_{\text{ext}} e_z$, it maps to the vector field $\operatorname{curl} \mathbf{h}_{\text{ext}} = (-\partial_y h_{\text{ext}}, \partial_x h_{\text{ext}})^\top$ whereas for a vector field \mathbf{A} the curl maps to the scalar field $\operatorname{curl} \mathbf{A} = \partial_x A_2 - \partial_y A_1$. For simplicity, we will use the same symbol, curl, for both

operators, as the context will make it clear which is meant. For any vector $a, b \in \mathbb{R}^2$ we formally set $a \times b = a_1 b_2 - a_2 b_1$. These definitions extend the spaces $\mathbf{H}(\text{curl})$ and $\mathbf{H}(\text{div})$ and their corresponding subspaces of zero traces to two dimensions in a straightforward way.

As mentioned above, the state of a superconductor is described by a minimizing pair $(u, \mathbf{A}) \in H^1 \times \mathbf{H}(\text{curl})$ of the GL free energy, i.e.,

$$E(u, \mathbf{A}) = \min_{(v, \mathbf{B}) \in H^1 \times \mathbf{H}(\text{curl})} E(v, \mathbf{B}).$$

The first Fréchet derivative of the energy functional

$$E'(u, \mathbf{A}) = (\partial_u E(u, \mathbf{A}), \partial_{\mathbf{A}} E(u, \mathbf{A})) : H^1 \times \mathbf{H}(\text{curl}) \rightarrow \mathbb{R}$$

is given by

$$(2.1a) \quad \langle \partial_u E(u, \mathbf{A}), v \rangle = \left(\frac{i}{\kappa} \nabla u + \mathbf{A}u, \frac{i}{\kappa} \nabla v + \mathbf{A}v \right) + ((|u|^2 - 1)u, v)$$

$$(2.1b) \quad \langle \partial_{\mathbf{A}} E(u, \mathbf{A}), \mathbf{B} \rangle = \left(|u|^2 \mathbf{A} + \frac{1}{\kappa} \text{Re}(i\bar{u}\nabla u), \mathbf{B} \right) + (\text{curl } \mathbf{A} - \mathbf{h}_{\text{ext}}, \text{curl } \mathbf{B})$$

for $(v, \mathbf{B}) \in H^1 \times \mathbf{H}(\text{curl})$. By the first-order optimality condition each minimizer (u, \mathbf{A}) , if it exists, needs to be a critical point of E , i.e., $\langle E'(u, \mathbf{A}), (v, \mathbf{B}) \rangle = 0$ for all $(v, \mathbf{B}) \in H^1 \times \mathbf{H}(\text{curl})$. This leads to the famous GL equations

$$\begin{aligned} & \left(\frac{i}{\kappa} \nabla u + \mathbf{A}u, \frac{i}{\kappa} \nabla v + \mathbf{A}v \right) + ((1 - |u|^2)u, v) = 0, \\ & \left(|u|^2 \mathbf{A} + \frac{1}{\kappa} \text{Re}(i\bar{u}\nabla u), \mathbf{B} \right) + (\text{curl } \mathbf{A} - \mathbf{h}_{\text{ext}}, \text{curl } \mathbf{B}) = 0, \end{aligned}$$

for all $(v, \mathbf{B}) \in H^1 \times \mathbf{H}(\text{curl})$. One proves that indeed there exists at least one minimizing pair. Nevertheless, any such minimizer cannot be unique due to so-called *gauge-transformations*. Let $(u, \mathbf{A}) \in H^1 \times \mathbf{H}(\text{curl})$ be arbitrary. For any $\varphi \in H^2$ we define the gauge transformation

$$G_\varphi : H^1 \times \mathbf{H}(\text{curl}) \rightarrow H^1 \times \mathbf{H}(\text{curl}), \quad G_\varphi(u, \mathbf{A}) = (e^{i\kappa\varphi}u, \mathbf{A} + \nabla\varphi).$$

It now holds that $E(u, \mathbf{A}) = E(G_\varphi(u, \mathbf{A}))$, i.e., the GL energy is invariant with respect to the gauge transformation, and hence a minimizer in $H^1 \times \mathbf{H}(\text{curl})$ cannot be unique. Although we have to deal with the symmetry under complex phase shifts the order parameter u , we can remove the gauge invariance with respect to the vector potential \mathbf{A} by restricting the minimization to divergence-free vector potentials. For this so-called *Coulomb gauge*, we employ the space

$$\mathbf{H}_0(\text{curl}, \text{div}^0) = \{\mathbf{B} \in \mathbf{H}_0(\text{curl}, \text{div}) \mid \text{div } \mathbf{B} = 0\}.$$

Next, we recall a result, which shows that at least one minimizer of the GL energy (1.1) exists. Moreover, the minimizer can be chosen from the space $\mathbf{H}_0(\text{curl}, \text{div}^0)$, making use of the gauge transformation, see [11, Lemma 2.6] and [13, Theorem 3.5].

PROPOSITION 2.1. *Let $\Omega \subset \mathbb{R}^d$, $d = 2, 3$ be a bounded, simply connected Lipschitz domain and $\mathbf{h}_{\text{ext}} \in L^2(\Omega, \mathbb{R}^{\ell_d})$. Then, the following statements hold.*

- (1) *Let $(u, \mathbf{A}) \in H^1 \times \mathbf{H}(\text{curl})$ be arbitrary. Then, there exists a function $\varphi \in H^1$ such that $(\tilde{u}, \tilde{\mathbf{A}}) := G_\varphi(u, \mathbf{A}) \in H^1 \times \mathbf{H}_0(\text{curl}, \text{div}^0)$ and $E(\tilde{u}, \tilde{\mathbf{A}}) = E(u, \mathbf{A})$.*

(2) *There exists at least one minimizer $(u, \mathbf{A}) \in H^1 \times \mathbf{H}_0(\text{curl}, \text{div}^0)$ of E .*

According to Proposition 2.1, it suffices to seek for a minimizing pair (u, \mathbf{A}) with a divergence-free vector potential. We enforce the vanishing divergence of the vector potential by a weak side constraint leading to the following optimization problem: Find $(u, \mathbf{A}) \in H^1 \times \mathbf{H}(\text{curl})$ such that

$$(2.2) \quad E(u, \mathbf{A}) = \min_{(v, \mathbf{B}) \in H^1 \times \mathbf{H}(\text{curl})} E(v, \mathbf{B}) \quad \text{subject to } (\mathbf{A}, \nabla v)_{\mathbf{L}^2} = 0 \text{ for all } v \in H^1.$$

Therefore, every such minimizer automatically belongs to $H^1 \times \mathbf{H}_0(\text{curl}, \text{div}^0)$, i.e., \mathbf{A} is divergence-free and satisfies the boundary condition $\mathbf{A} \cdot \nu = 0$ on $\partial\Omega$.

Due to the constraint of vanishing divergence in \mathbf{A} the first-order conditions read as a saddle point problem involving a Lagrange multiplier to enforce the side constraint. A minimizer $(u, \mathbf{A}) \in H^1 \times \mathbf{H}(\text{curl})$ of (2.2) is also characterized by the solution triple $(u, \mathbf{A}, \lambda) \in H^1 \times \mathbf{H}(\text{curl}) \times H^1$ that solves the saddle point problem

$$(2.3) \quad \begin{aligned} \langle E'(u, \mathbf{A}), (v, \mathbf{B}) \rangle + (\nabla \lambda, \mathbf{B})_{\mathbf{L}^2} &= 0 \quad \text{for all } (v, \mathbf{B}) \in H^1 \times \mathbf{H}(\text{curl}), \\ (\mathbf{A}, \nabla w)_{\mathbf{L}^2} &= 0 \quad \text{for all } w \in H^1. \end{aligned}$$

Note that $\lambda \in H^1$ denotes the associated Lagrange multiplier. Conversely, every critical point solving (2.3) corresponds to either a local minimum or a saddle point of the energy functional E . Loosely speaking, such a critical point with a low energy is a good candidate for the global minimizer if we verify the second-order optimality condition for a local minimum. That is, we have to exclude the possibility of a saddle point.

2.2. A reduced Ginzburg–Landau model with given \mathbf{A} . To obtain first impressions on the applicability and potential of our approach, in this work we will also consider a reduced model of (1.1) where we assume \mathbf{A} to be given by the relation $\text{curl } \mathbf{A} = \mathbf{h}_{\text{ext}}$. The corresponding modified energy then reads

$$(2.4) \quad E(u) := \frac{1}{2} \int_{\Omega} \left| \frac{i}{\kappa} \nabla u + \mathbf{A}u \right|^2 + \frac{1}{2} (1 - |u|^2)^2 \, dx,$$

and is only minimized for u . This has the advantage of fewer degrees of freedom while it still provides deep insight into many challenges.

2.3. Numerical computation of Ginzburg–Landau minimizers. In the following section, we briefly discuss two approaches which are purely neural network-driven or finite element-driven, and focus mainly on their advantages and disadvantages. A thorough explanation is postponed to Section 3 and to Section 4 for the neural network and the finite element solver, respectively. The main contribution of our work is to study suitable combinations of both approaches, which we will develop at the end of this section. The corresponding numerical experiments are presented in Section 5.

Neural network approximation. A modern approach is the use of *neural networks* (NNs). They are able to exploit modern architectures within *graphics processing units* (GPUs) and their compute-heavy training is easily parallelizable. Several approaches for NNs with supervised and unsupervised learning have been developed over recent years ranging from data-driven models to physics-based NNs. For most of these approaches, the strategy to train the NN is to minimize a prescribed loss functional. In a supervised setting, this could be a certain distance between the NN

output and the labels of the training data. However, our setting fits much better to an unsupervised learning strategy, since the GL energy (1.1) gives rise to a natural loss function. This could be seen as a variant of the deep Ritz method. Another huge advantage of the usage of NNs is that there are many powerful optimization algorithms freely available in standard Python libraries such as `PyTorch`, which allow for the training of very large models.

However, finding a suitable architecture for a network as well as achieving a successful training process is rather challenging. More precisely, the choice of the precise network architecture (including many tunable parameters) itself is a high-dimensional optimization problem. Further, the resulting minimization problem in the training is generally non-convex. Thus, it is very difficult to reliably find a good network that is not just an unsatisfactory local minimum. Due to the large number of degrees of freedom, randomness at different stages of the training process is essential in order to obtain reasonable gradient descent methods. Examples are the initialization of the weights and the use of stochastic gradient descent methods. While randomness is computationally advantageous, it also means that from a theoretical perspective deterministic and reliable error estimates are typically not available.

Finite element approximations. A more classical approach to compute minimizers such as in (1.1) consists in the use of finite element (FE) methods. Here, one replaces the infinite-dimensional spaces by finite-dimensional approximation spaces, e.g., spaces of piecewise polynomials. Then one ends up with a finite-dimensional optimization problem, for which well-known methods can be used. One advantage is that the resulting matrices are in general sparse and thus pave the way for efficient computations in the minimization procedure. In addition, it is possible to show rigorous approximation rates and resolution conditions in terms of the mesh size h and the parameter κ as mentioned in the introduction. Further, computing a-posteriori the smallest eigenvalues of $E''(u)$ can be used to avoid ending up in a saddle point. In summary, this provides more reliability in the computed approximations compared to the pure NN approach.

However, there are also some disadvantages of this approach. Compared to the NN implementation, leveraging the computational potential of GPUs in a FE solver is not trivial, and thus much harder to compete with the resulting computational speed-up, even using an MPI (Message Passing Interface) implementation as done in our code. In our experiments, we also observed a strong sensitivity to the chosen initial values leading to completely different discrete minimizers on different energy levels. Unfortunately, the high cost for a single minimization run prevents us from exploring the initial value landscape in a systematic manner.

A neural network enhanced two-step approach (hybrid-approach). Our idea is to combine the advantages of both above-mentioned methods in a so-called *hybrid approach*. More precisely, we use the NN to compute a first approximation of the minimizer, where we can exploit the great parallelization on GPUs. We then interpolate the minimizer into the FE space and use it as an initial value for the iterative FE solver. Thus, one could interpret the NN outcome as a prediction for the exact FE minimizer. We thus have the advantage of replacing the heuristic choices of initial values with the more sophisticated surrogates of the NN while computing the final minimizer with the reliability of the FE solver. In addition, we observe later on that it is beneficial for the NN to compute the minimizer on a whole interval of admissible κ values. The aim is then that a good choice of a small number of initial values excludes several local minimizers. Overall, we pursue the following two

approaches:

- We perform a training of the NN with a problem-specific objective to compute approximations of GL energy minimizers. We train the network on a range of κ and interpolate it for a fixed κ into the FE space to compute energies that are comparable to the FE solver. Since the resulting network can be used on the full range of κ , one can then decide a-priori for how many values of κ the network shall be used in order to budget the time for the training accordingly. The evaluation of the NN after the training to generate solutions for different κ values is negligible regarding the runtime when compared to a FE solver.
- We use the same procedure as in the previous point, but use this interpolation as an initial value for the FE solver, to possibly achieve states of even lower energy. This enhances the FE solver with the NN and thus also indirectly allows the FE solver to profit from GPU computations without any adaptation on the FE side of the code.

We elaborate on the precise strategies in our numerical examples in Section 5.

3. Neural networks. In this section, we introduce the NN approach that we use to solve the GL energy minimization problem (2.2). As mentioned before, this approach can be used as a stand-alone method to compute approximations of GL energy minimizers or as the first part of the hybrid-approach to compute an initial guess for the second level conjugate Sobolev gradient iteration. The approach is applied to the full problem (1.1), as well as the reduced problem (2.4).

Before introducing the architecture and learning procedure of the NN for the GL problem, we first recall basic machine learning preliminaries for the sake of completeness. For further reading on the topic, we recommend [18].

3.1. Machine learning basics.

3.1.1. MLPs, training, optimizers, and more. One of the simplest NN architectures is realized by *multi-layer perceptron (MLP)*. MLPs are defined as a function $F = f^L \circ g^{L-1} \circ \dots \circ g^0$, where $g^j = \phi^j \circ f^j$. Here, f^j are affine-linear maps and ϕ^j are non-linear functions that are applied entry-wise. The ϕ^j are called *activation functions*. They raise the NN from being just an affine-linear map to a more expressive function. Note that NNs are not limited to said choice of the functions g^j . In fact, we use a different scheme in this work, see later. The individual maps g^j are also referred to as the *layers* of the NN, and the output dimension is typically called the *width* of the layer. Often, several layers with a specific structure are grouped together to form *blocks*. One very popular example of this are *residual layers/blocks*. A residual block can, for example, consist of a few MLP layers, the output of which is called the *residual branch* and is added to the input, hence the name residual. More formally, this means that instead of realizing $x \mapsto f_{\text{NN}}(x)$, where f_{NN} is the output of the MLP layers, the NN block realizes $\text{resBlock}(x) = (I + f_{\text{NN}})(x)$, where I is the identity map.

An MLP is parametrized by the matrices W^j (*weights*) and vectors b^j (*biases*) that define the affine-linear maps $f^j = W^j x + b^j$. A simple but typical machine learning task could now be to interpolate a set of points (*samples*) and values (*labels/targets*) with the NN. This set is then referred to as the *training dataset*. In order to interpolate the training samples, a loss functional, e.g., a norm or metric in the space of the labels, is evaluated. This loss is then minimized over the space of network parameters in an iterative process called *training*. To do this, one uses gradient descent like methods. The step size in these methods is referred to as the *learning rate*. The learning rate is often not constant over the gradient descent steps, but is scheduled. One

way this could look like is a gradual increase of the learning rate in the beginning (*warm-up*) [19], followed by a decay of the learning rate afterwards. Many of the *optimizers*, i.e., the algorithms used for the minimization task in machine learning, are so-called *momentum methods*. While a simple gradient descent gets stuck in a local minimum, momentum methods can overcome local minima. One very popular optimizer is *Adam* [21], which stands for “adaptive moment estimation”. A variant of Adam is *AdamW*. Here, the *weight decay*, which is a penalty term that penalizes large model weights, is decoupled from the rest of the optimization. AdamW is a default choice in machine learning, which is why we utilize it in this work. Another optimizer is *Muon* [20], which has shown good results in the training of large language models [22]. Due to said success, we use Muon in this work to evaluate whether it is a viable option for our setting as well. Muon is short for “momentum orthogonalized by Newton-Schulz” and essentially orthogonalizes the updates to the weight matrices, which can be beneficial. However, Muon can only be used on 2D parameters, which in our case only includes weight matrices. Other parameters still have to be trained with AdamW. For some of the networks trained, we used only AdamW, others were trained with Muon (with the adaptation from [22]) and AdamW in combination. When using both optimizers, AdamW was used for the scaling parameter of the residual branch (see below for further details). Overall, we did not find a clear advantage of one optimizer over the other. Nonetheless, we include both in our experiments for the sake of completeness.

Batches and epochs. Typically, the training dataset is processed in *batches*, which are elements in a random partitioning of the dataset. They are fed into the NN one after another and on each of the batches a training step is performed. After all batches have been processed, i.e., the entire dataset has been seen, one says that an *epoch* is completed. The batches are then re-drawn from the dataset and a new epoch begins. During training, one often computes *validation* metrics to check how good the NN performs outside of the training dataset.

The set of parameters defining the architecture and training strategy of a NN but are not trainable are called *hyperparameters* of the NN. These include, for example, the depth, width, and learning rate.

Unsupervised learning. Previously, we mentioned a training dataset consisting of samples and labels or targets. This style of training an NN is called *supervised learning*. But there is also *unsupervised learning*, which refers to the situation where no labels or targets are known. Instead, there are only samples without corresponding labels. The quality of the approximation of the target values is simply encoded in the loss functional. One example for such unsupervised learning tasks are *physics-informed neural networks (PINNs)* [26], where an NN is trained to solve a differential equation, by approximating the solution. The loss functional for a PINN is the mean squared error of the point-wise residual of the differential equation that, to some degree, measures how close the NN is to the solution without requiring the knowledge of said solution. The derivatives of the NN with respect to the input, that are required to compute the residual, can be computed using automatic differentiation. Another machine learning method to solve differential equations is the *deep Ritz method* [17]. Instead of the strong formulation, a weak formulation is used here. This weak formulation is stated as a minimization problem, which then is a natural problem formulation for the NN training. For the application of machine learning to the GLE, we follow an unsupervised method similar to the deep Ritz method.

3.2. Machine learning for the Ginzburg–Landau problem.

Loss functions. As indicated above, we use unsupervised learning for the GL problem. A major reason for this is that we do not have a dataset with target values. This means that we have to encode the metric of how good the NN performs in the loss function. For the GL problem, a canonical idea of a loss functional is the GL free energy. However, we also want to include the material parameter κ in the samples such that the trained model can be used for many runs of the solver. In order to do this, we integrate over a range of κ -values and also scale with the corresponding κ . In the discrete setting, this weighted integral reduces to a weighted sum.

The main motivation of this scaling is the asymptotic energy estimate in [28, Theorem 8.1], which (for constant \mathbf{h}_{ext}) yields $E(u, \mathbf{A}, \kappa) \sim \frac{\log \kappa}{\kappa}$ as $\kappa \rightarrow \infty$. However, since we do not treat arbitrarily large κ and we deviate from the constant magnetic field, we simply neglect the logarithmic term, but keep the linear term to compensate vanishing energies for large κ in the training. An additional benefit of this is that this gives a larger weight to samples corresponding to larger value of κ , for which the problems tends to be harder to solve. The loss functional for the GL problem reads

$$\mathcal{E}(u, \mathbf{A}) := \int_{\kappa_{\min}}^{\kappa_{\max}} \kappa \hat{E}(u, \mathbf{A}, \kappa) \, d\kappa,$$

where

$$\hat{E}(u, \mathbf{A}, \kappa) := \frac{1}{2} \int_{\Omega} \left| \frac{i}{\kappa} \nabla u + \mathbf{A}u \right|^2 + \frac{1}{2} (1 - |u|^2)^2 + |\operatorname{curl} \mathbf{A} - \mathbf{h}_{\text{ext}}|^2 + |\operatorname{div} \mathbf{A}|^2 \, dx.$$

For the reduced GL problem (2.4), we simply plug in the considered \mathbf{A} . We also experimented with models trained for only a single κ value, where we used $\hat{E}(u, \mathbf{A}, \kappa)$ as the training loss. However, one of the findings of the preliminary experiments is that – especially in the case of solving only for the order parameter – training for a range of admissible κ seems to yield better results for higher values of κ than just training for the specific value. Particularly, increasing κ_{\max} beyond the actual target range to, e.g., $\kappa_{\max} = 130$, improved the results for $\kappa = 100$.

The procedure of learning the dependence of the solution on κ can be interpreted for the GL problem as learning the parameter-to-solution map

$$S: [\kappa_{\min}, \kappa_{\max}] \rightarrow H^1(\Omega, \mathbb{C} \times \mathbb{R}^d), \quad \kappa \mapsto (u_{\kappa}, \mathbf{A}_{\kappa}) = \operatorname{argmin}(E(u, \mathbf{A}, \kappa)).$$

A similar perspective can be chosen for the reduced problem.

Neural network architecture. In contrast to the previously discussed simple architectures of NNs via MLPs, we use more sophisticated modern architectures, which turns out to be necessary and that we describe in the following. The network consists of a variable number of blocks with a variable width. Each block has a gated structure [8]. Gating in this context means entry-wise multiplication of two branches, one of which is activated and called *gate*. In particular, we make use of *SwiGLU* layers [32], which have proven to be useful in the *feed-forward* blocks of *transformers*, which are network architectures that are used, e.g., in large language models. Feed-forward layers in this context can, for example, be MLP layers or, as just mentioned, SwiGLU layers. The main idea in the SwiGLU scheme is, that the input of the layer is split up into two linearly mapped branches. Then, one of the branches is activated and the result is multiplied entry-wise with the second branch. Afterwards, the result is once again mapped via a linear map, see (3.1). We also propose a slight variation with a double activation scheme (which we refer to

as *DAGLU*). Here, an additional linear layer followed by an activation function is placed at the beginning of the block. In exchange, the final linear layer on the output is removed. This scheme can be combined with GELU or SiLU activations. Such structures have turned out to be advantageous in preliminary tests, which is why we included them in our constructions. The SwiGLU and DAGLU block formulas read

$$(3.1) \quad \text{SwiGLU}(x, W_1, W_2, V) = W_1(\text{SiLU}(W_2x) \otimes Vx),$$

and

$$\text{DAGLU}(x, W_1, W_2, V, \alpha) = \alpha(W_2h) \otimes Vh, \quad h = \alpha(W_1x), \quad \alpha \in \{\text{SiLU}, \text{GELU}\}.$$

where \otimes is the entry-wise multiplication of vectors and W_1, W_2 and V are appropriately sized matrices.

The formulas for activation functions used above read

$$\text{SiLU}(x) = \frac{x}{1 + \exp(-x)} \quad \text{and} \quad \text{GELU}(x) = x\Phi(x),$$

where Φ is the cumulative Gaussian distribution function.

The models we use have a residual structure in conjunction with the SwiGLU and DAGLU schemes to form blocks. The residual branch is scaled by a factor that can also be learned or fixed. We mostly use a learned entry-wise scaling vector [33] with values in $(0, 1]$, initialized by $1/(2L)$, where L is the number of blocks in the network. The small and layer-dependent initialization helps with stability in deeper networks without using normalization layers and is similar in concept to [35]. Using normalization layers in the NN proved detrimental for convergence in our preliminary testing.

The resulting formula for the SwiGLU block including the residual structure is

$$\text{RSwiGLU}(x, W_1, W_2, V, \gamma) = x + \gamma \otimes \text{SwiGLU}(x, W_1, W_2, V),$$

and γ is a vector of scaling factors and, as before, \otimes refers to the entry-wise multiplication of vectors. The DAGLU block including the residual structure reads

$$\text{RDAGLU}(x, W_1, W_2, V, \alpha, \gamma) = x + \gamma \otimes \text{DAGLU}(x, W_1, W_2, V, \alpha),$$

where again, γ is a vector of scaling factors and \otimes is the entry-wise multiplication of vectors. For a visual representation of the SwiGLU and DAGLU block scheme, see Figure 1.

Hardware. The training of the NNs used in the experiments is done on the HoreKa Tier 2 High Performance Computing system at KIT. The training uses four nodes with four NVIDIA A100 GPUs per node for a total of 16 GPUs. For the tuning and preliminary experimentation, we use the HAICORE@KIT partition of HoreKa. For the interpolation of the NNs, a GPU should be used. We leverage a workstation with a single NVIDIA A100 GPU for this purpose, which proved to be able to compute the forward passes reasonably fast.

Learning rate. The scheduling of the learning rate during the training is divided into several phases. First, we do a short gradual warm-up [19] of the learning rate to prevent instabilities. Here, the learning rate starts at a smaller value and is then increased over a set number of training steps. Second, a gradual decay of the learning rate is done using either an exponential decay or a cosine decay. The latter is a

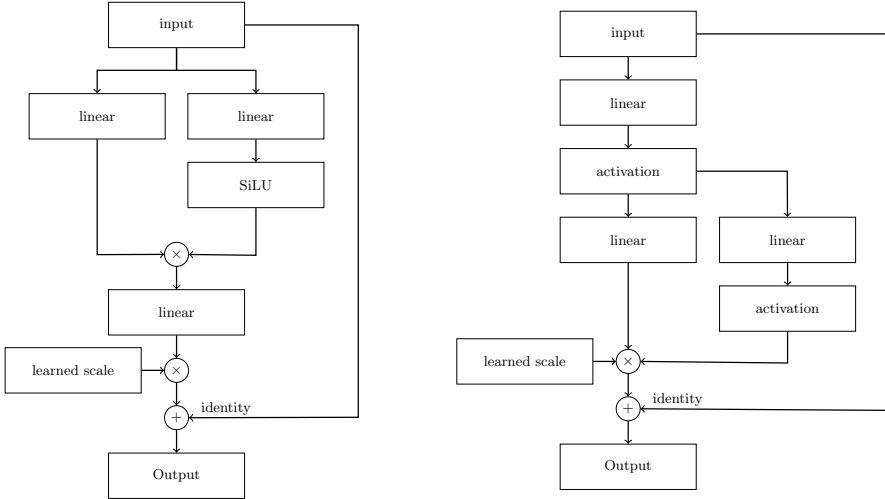


Fig. 1: Block schematic of a SwiGLU block (left) and a DAGLU block (right).

decay which follows a cosine in its shape and decays the learning rate down to a set minimum value. In the case of the cosine decay, we do an additional linear decay of the learning rate for a number of steps after the cosine scheduling ends. For the exponential scheduling, we skip this part in our experiments, i.e., the training was stopped before the linear decay phase. For the exact hyperparameters used for each of the models in this work, we refer to Table 1 and the configuration files in the source code that is available upon request. There is no particular reason as to why which strategy was used for which model. There is no clear preference of one of the scheduling strategies over the other in our results and we include both strategies for completeness.

Datasets. For the training of the networks including the material parameter κ , the dataset consists of random samples in $\Omega \times [\kappa_{\min}, \kappa_{\max}]$. The samples are drawn in bulk every epoch. These samples are then processed in batches until each sample has been used once, then the samples are re-drawn with a new seed.

4. Finite element approximation and gradient descent method. In order to obtain a FE formulation of the GL problem, we simply restrict the minimization in (2.2) to some finite-dimensional subspace. For our purposes, we choose continuous finite elements of order $p \in \mathbb{N}$ for the real and imaginary part of the order parameter u , denoted by V_h , and *Nédélec elements* of first kind and order m , denoted by \mathcal{R}_h . Nédélec elements are the natural choice for our approach, since they only enforce continuity in the tangential direction on each element. Further, since they are not a subspace of \mathbf{H}^1 , they can still capture the correct solution behavior even in the case of non-convex domains. In order to enforce the divergence constraint, we have to define the weak divergence-free condition. To this end let \tilde{V}_h denote the space of continuous elements of order m , and define the discrete divergence as

$$(4.1) \quad \operatorname{div}_h \mathbf{B}_h = 0 \iff (\mathbf{B}_h, \nabla \varphi_h) = 0 \text{ for all } \varphi_h \in \tilde{V}_h.$$

Let us note that this does not only enforce that the weak divergence vanishes, but in addition also that the normal trace is zero in a weak sense. Then, the discrete

minimization problem reads: seek $(u_h, \mathbf{A}_h) \in V_h \times \mathcal{R}_h$ such that

$$(4.2) \quad E(u_h, \mathbf{A}_h) = \min_{(v_h, \mathbf{B}_h) \in V_h \times \mathcal{R}_h} E(v_h, \mathbf{B}_h) \quad \text{subject to } \operatorname{div}_h \mathbf{A}_h = 0.$$

By changing the spaces, the minimizer of (4.2) satisfies the discrete analogue of (2.3). Further, we define $\mathcal{R}_h^0 = \{\mathbf{B}_h \in \mathcal{R}_h \mid \operatorname{div}_h \mathbf{B}_h = 0\}$ as the discrete counterpart of $\mathbf{H}_0(\operatorname{curl}, \operatorname{div}^0)$ such that (4.2) can be formulated as

$$E(u_h, \mathbf{A}_h) = \min_{(v_h, \mathbf{B}_h) \in V_h \times \mathcal{R}_h^0} E(v_h, \mathbf{B}_h).$$

Let us emphasize that this FE approach is non conforming since a discrete divergence free function in (4.1) is in general not divergence free, i.e., $\mathcal{R}_h^0 \not\subseteq \mathbf{H}_0(\operatorname{curl}, \operatorname{div}^0)$, but it only holds the weaker inclusion $\mathcal{R}_h^0 \subset \mathbf{H}(\operatorname{curl})$.

Next, we propose a special iterative solver to compute a reliable approximation of a discrete GL energy minimizer. Since we only apply it to the discretized energy (4.2), we derive the algorithm only for the discrete case, even though this can be done in greater generality. To keep the notation simple, we also refrain from adding lower indices h to every variable. We emphasize that this method was recently proposed in [7] for the case reduced case (2.4), and our method is a generalization of this approach to the full problem (1.1).

Starting from an initial guess (u^0, \mathbf{A}^0) in the FE space (which for example can be obtained by some NN), the iteration takes the form of a descent method given by

$$(u^{k+1}, \mathbf{A}^{k+1}) = (u^k, \mathbf{A}^k) + \tau_k (d^k, \mathbf{D}^k), \quad k = 0, 1, 2, \dots,$$

where $(d^k, \mathbf{D}^k) \in V_h \times \mathcal{R}_h^0 \subset H^1 \times \mathbf{H}(\operatorname{curl})$ is the descent direction that we specify in the following. For this purpose, we equip the tensor space $X^k := V_h \times \mathcal{R}_h^0$ with the k -dependent metric

$$((v, \mathbf{B}), (w, \mathbf{C}))_{X^k} := a_k(v, w) + b_k(\mathbf{B}, \mathbf{C}), \quad (v, \mathbf{B}), (w, \mathbf{C}) \in X^k,$$

where

$$\begin{aligned} a_k(v, w) &:= \left(\frac{1}{\kappa} \nabla v + \mathbf{A}^k v, \frac{1}{\kappa} \nabla w + \mathbf{A}^k w \right)_{L^2} + \left((\beta + |u^k|^2 + |\mathbf{A}^k|^2) v, w \right)_{L^2} \\ b_k(\mathbf{B}, \mathbf{C}) &:= (\operatorname{curl} \mathbf{B}, \operatorname{curl} \mathbf{C})_{L^2} + \left((\beta + |u^k|^2) \mathbf{B}, \mathbf{C} \right)_{L^2} \end{aligned}$$

for some stabilization parameter $\beta > 0$. In particular, $(X^k)_{k \in \mathbb{N}}$ forms a sequence of Hilbert spaces defined by the same underlying vector space, but with a varying metric in each iteration which is indicated by the super index. Note that if β is sufficiently large (independent of k), then $a_k(\cdot, \cdot)$ and $b_k(\cdot, \cdot)$ are coercive and bounded on V_h and \mathcal{R}_h^0 , respectively. Therefore, $(\cdot, \cdot)_{X^k}$ is indeed an inner product on the vector space X^k . However, the subsequent derivation of the conjugate Sobolev gradient method also applies for other choices as long as $(X^k, (\cdot, \cdot)_{X^k})$ is a Hilbert space and $X^k \subset H^1 \times \mathbf{H}(\operatorname{curl})$.

For each X^k , we can formally define the Sobolev gradient $\nabla_{X^k} E(u^k, \mathbf{A}^k)$ of the GL energy functional at the current iterate $(u^k, \mathbf{A}^k) \in X^k$ by

$$(4.3) \quad (\nabla_{X^k} E(u^k, \mathbf{A}^k), (v, \mathbf{B}))_{X^k} = \langle E'(u^k, \mathbf{A}^k), (v, \mathbf{B}) \rangle \quad \text{for all } v \in V_h, \mathbf{B} \in \mathcal{R}_h^0,$$

where $E'(u^k, \mathbf{A}^k) \in (V_h \times \mathcal{R}_h^0)^*$ denotes the Frechét derivative and $\langle \cdot, \cdot \rangle$ is the dual pairing between $H^1 \times \mathbf{H}(\operatorname{curl})$ and its dual. Roughly speaking, the metric $(\cdot, \cdot)_{X^k}$ is

chosen in a way such that the Sobolev gradient $\nabla_{X^k} E(u, \mathbf{A})$ is close to the identity map. Recalling the Frechét derivatives in the directions of u and \mathbf{A} given in (2.1a) and (2.1b), respectively, a straightforward calculation shows that the Sobolev gradient (4.3) is given by

$$\nabla_{X^k} E(u^k, \mathbf{A}^k) = (u^k, \mathbf{A}^k) - (\delta^k, \Delta^k)$$

where $\delta^k \in V_h$ is the solution of the elliptic problem

$$a_k(\delta^k, v) = ((1 + \beta + |\mathbf{A}^k|^2)u^k, v)_{L^2} \quad \text{for all } v \in V_h$$

and $\Delta^k \in \mathcal{R}_h^0$ is the solution of the elliptic problem

$$b_k(\Delta^k, \mathbf{B}) = (\beta \mathbf{A}^k - \frac{1}{\kappa} \operatorname{Re}(i\bar{u}^k \nabla u^k), \mathbf{B})_{\mathbf{L}^2} + (\mathbf{h}_{\text{ext}}, \operatorname{curl} \mathbf{B})_{\mathbf{L}^2} \quad \text{for all } \mathbf{B} \in \mathcal{R}_h^0.$$

The descent direction is now defined in the spirit of a nonlinear conjugate Sobolev gradient descent method by

$$(d^k, \mathbf{D}^k) := \begin{cases} -\nabla_{X^k} E(u^k, \mathbf{A}^k), & k = 0 \\ -\nabla_{X^k} E(u^k, \mathbf{A}^k) + \gamma_k (d^{k-1}, \mathbf{D}^{k-1}), & k \geq 1 \end{cases}$$

where γ_k is the Polak-Ribière dissipation parameter, cf. [25], defined by

$$\gamma_k = \max \left\{ 0, \frac{(\nabla_{X^k} E(u^k, \mathbf{A}^k), \nabla_{X^k} E(u^k, \mathbf{A}^k) - \nabla_{X^{k-1}} E(u^{k-1}, \mathbf{A}^{k-1}))_{X^k}}{(\nabla_{X^{k-1}} E(u^{k-1}, \mathbf{A}^{k-1}), \nabla_{X^{k-1}} E(u^{k-1}, \mathbf{A}^{k-1}))_{X^{k-1}}} \right\}.$$

We further orthogonalize the search direction for \mathbf{A} in the sense of (4.1) such that $\operatorname{div}_h \mathbf{D}^k = 0$, which ensures the discrete divergence condition for the next approximation \mathbf{A}^k . Along the descent direction, we can then look for the optimal step size τ_k that minimizes the energy

$$\tau_k = \operatorname{argmin}_{\tau > 0} E((u^k, \mathbf{A}^k) + \tau (d^k, \mathbf{D}^k)).$$

Note that $E((u^k, \mathbf{A}^k) + \tau (d^k, \mathbf{D}^k))$ is a fourth order polynomial in τ , which we denote by p_k . Then, τ_k can be computed efficiently by a line search or we compute the roots of p'_k directly.

We can now formulate the final algorithm for the nonlinear conjugate Sobolev gradient descent method.

Algorithm 4.1 (nonlinear conjugate Sobolev gradient descent method). Let initial values $u^0 \in V_h$ and $\mathbf{A}^0 \in \mathcal{R}_h^0$ with $u^0 \neq 0$ and a tolerance $\text{tol} > 0$ be given. Further, set $d^0 = 0$ and $\mathbf{D}^0 = 0$. Then for $k = 0, 1, 2, \dots$ perform the following steps:

1. **Solve** for $\delta^k \in V_h$ such that

$$a_k(\delta^k, v) = ((1 + \beta + |\mathbf{A}^k|^2)u^k, v)_{L^2} \quad \text{for all } v \in V_h$$

and for $\Delta^k \in \mathcal{R}_h$ such that

$$b_k(\Delta^k, \mathbf{B}) = (\beta \mathbf{A}^k - \frac{1}{\kappa} \operatorname{Re}(i\bar{u}^k \nabla u^k), \mathbf{B})_{\mathbf{L}^2} + (\mathbf{h}_{\text{ext}}, \operatorname{curl} \mathbf{B})_{\mathbf{L}^2} \quad \text{for all } \mathbf{B} \in \mathcal{R}_h.$$

2. **Update** the descent directions

$$d^k = \delta^k - u^k + \gamma_k d^{k-1}, \quad \mathbf{D}^k = \Delta^k - \mathbf{A}^k + \gamma_k \mathbf{D}^{k-1}$$

with the Polak-Ribière parameter $\gamma_k = \max\{0, \frac{N_k}{B_k}\}$ with

$$B_k = ((\delta^k - u^k, \Delta^k - \mathbf{A}^k), (\delta^k - u^k, \Delta^k - \mathbf{A}^k))_{X^{k-1}}$$

$$N_k = ((\delta^k - u^k, \Delta^k - \mathbf{A}^k), (\delta^k - u^k - \delta^{k-1} + u^{k-1}, \Delta^k - \mathbf{A}^k - \Delta^{k-1} + \mathbf{A}^{k-1}))_{X^k}$$

and orthogonalize such that $\operatorname{div}_h \mathbf{D}^k = 0$.

3. **Determine** the step size τ_k such that

$$\tau_k = \operatorname{argmin}_{\tau > 0} E((u^k, \mathbf{A}^k) + \tau(d^k, \mathbf{D}^k))$$

via a line search (e.g., golden search) or direct computation of the roots of p'_k .

4. **Iterate** according to $(u^{k+1}, \mathbf{A}^{k+1}) = (u^k, \mathbf{A}^k) + \tau_k(d^k, \mathbf{D}^k)$.

5. **Terminate** if $E(u^k, \mathbf{D}^k) - E(u^{k+1}, \mathbf{D}^{k+1}) < \text{tol}$.

In practice, it is sufficient to set $\beta = 0$. The reason is that the iterates (u^k, \mathbf{A}^k) typically contribute enough mass through the second terms in the bilinear forms $a_k(\cdot, \cdot)$ and $b_k(\cdot, \cdot)$ to ensure coercivity on V_h and \mathcal{R}_h^0 , respectively.

5. Numerical results. In this section, we present the numerical results obtained using the NN approach implemented in PyTorch [4], as well as the hybrid approach that combines the NN with the FE iterative solver implemented in DOLFINx [5]. We compare these results to those computed by the FE iterative solver using empirically chosen initial values to emphasize the advantages of the hybrid approach. We investigate various GL settings and training strategies for NNs. First, we consider NNs that are intensively trained for the reduced and full GL models. These networks lead to minimizers on lower energy levels than the ones obtained with the classical FE iterative approach with heuristic initial values. These models also provide a set of architectures and training parameters used in subsequent experiments involving faster-trained networks for GL models defined on different geometric domains, i.e., the unit square and the L-shaped domain. Table 1 depicts the most relevant parameters for the architecture and training strategy of the NN models used in all the experiments. All our models consist of 40 blocks of varying type with a width of 512 neurons resulting in about 31.5M trainable parameters. The main differences between the models lie in the chosen block type as well as the optimizer used for training. Regarding the learning rates, `target_lr_main` is the learning rate used for the weights, while `target_lr_scaling` is the learning rate used for the scaling vector for the residual branch. Further details on the hyperparameters can be found in the corresponding configuration files. These files are provided with the source code for the entire implementation, available at:

<https://github.com/BenjaminDoerich/glenn>

For the hybrid approach and for computing reference approximations, we exploit the FE iterative solver with quadratic Lagrange FE in the order parameter u and Nédélec elements of first kind and lowest order for the vector potential \mathbf{A} . On the unit square domain, a quasi-uniform mesh is used with 608,842 real (or 304,421 complex) degrees of freedom for the order parameter and 228,059 degrees of freedom for the

	GL-R1	GL-R2	GL-F1	GL-F2	GL-L*
depth (blocks)	40	40	40	40	40
width	512	512	512	512	512
trainable param.	31.5M	31.5M	31.5M	31.5M	31.5M
block type	SwiGLU	SwiGLU	SwiGLU	DAGLU	SwiGLU
activation	SiLU	SiLU	SiLU	SiLU	SiLU
precision	32	64	32	64	32
batch size	2^{14}	2^{15}	2^{15}	2^{15}	2^{15}
steps per epoch	2^{13}	2^{13}	2^{13}	2^{13}	2^{13}
total epochs	20	18	19	6	4
target lr main	0.001	0.003	0.003	0.003	0.01
target lr scaling	0.001	0.0003	0.0003	0.0003	0.001
decay scheduler	cosine	exp.	exp.	exp.	exp.
optimizer	AdamW	AdamW	Muon	Muon	Muon
κ_{\max}	110	130	130	130	130

Table 1: Architecture and training parameters for the NN models.

vector potential, which amounts to a total of 836,901 real degrees of freedom. For the L-shaped domain we have 457,306 real (or 228,653 complex) degrees of freedom for the order parameter and 170,233 degrees of freedom for the vector potential. This amounts to a total of 628,539 real degrees of freedom. Depending on the precise setting, one FE step takes around 5–7 seconds, leading to 80–120 minutes per 1000 iterations. The iterative solver for the classical and hybrid approach terminates once a tolerance of $\text{tol} = 10^{-12}$ is reached with respect to the energy increments.

As mentioned above, to quantify the quality of the NN and hybrid approach we compute reference approximations using the classical FE iterative solver with empirically chosen initial values. These are denoted by φ_j , $j = 1, 2, 3, 4, 5$ and generated for $\alpha = \frac{1}{\sqrt{2}}(1 + i)$ through the functions

$$\begin{aligned} \psi_1(x) &= \alpha e^{-|x|^2}, & \psi_2(x) &= \frac{1}{\sqrt{\pi}} \left(\frac{2}{3}(x_1 + ix_2) + \frac{1}{2} \right) e^{-|x|^2}, & \psi_3(x) &= \alpha e^{10i|x|^2}, \\ \psi_4(x) &= \alpha(x_1 + ix_2)e^{10i|x|^2}, & \psi_5(x) &= \alpha. \end{aligned}$$

With the transformation $\chi : (0, 1)^2 \rightarrow (-1, 1)^2$, $\chi(x) = (2x_1 - 1, 2x_2 - 1)$, we set the initial values for the order parameter on the unit square to interpolation into the FE space of $u^0 = \varphi_j$, $\varphi_j = \psi_j \circ \chi$ for $j = 1, 2, 3, 4, 5$. For the L-shaped domain we restrict these functions accordingly. In each GL setting, this leads to five reference approximations from which the one with the lowest energy level is the best candidate for the global minimizer obtained with the classical FE iterative solver.

5.1. Neural networks for the reduced GL model. For the first experimental setup, we consider the reduced model (2.4) on the unit square $\Omega = (0, 1)^2$. The magnetic vector potential \mathbf{A} is given by

$$\mathbf{A}(x) = \sqrt{2} \begin{pmatrix} \sin(\pi x_1) \cos(\pi x_2) \\ -\cos(\pi x_1) \sin(\pi x_2) \end{pmatrix}, \quad x = (x_1, x_2) \in \Omega.$$

Note that $\operatorname{div} \mathbf{A} = 0$ and that $\operatorname{curl} A = \mathbf{h}_{\text{ext}}$ with $\mathbf{h}_{\text{ext}}(x) = 2\sqrt{2}\pi \sin(\pi x_1) \sin(\pi x_2)$. The GL parameter is set to the values $\kappa = 10, 25, 50, 75, 100$. We emphasize again that the training of the NN does not particularly include these values, but rather provides approximations for a whole range of values of $\kappa \in [\kappa_{\min}, \kappa_{\max}]$.

κ	10	25	50	75	100
φ_1	0.10459064	0.08611814	0.06928056	0.03607700	0.04583871
φ_2	0.11497944	0.07431313	0.06231434	0.04038106	0.04301528
φ_3	0.10459064	0.06483810	0.04141371	0.03611419	0.03091586
φ_4	0.10459064	0.06483810	0.04141371	0.03448296	0.03091576
φ_5	0.10459064	0.08386487	0.06876863	0.04039311	0.04045206
GL-R1 (NN)	0.10462698	0.06523432	0.04215652	0.03212462	0.02785582
GL-R2 (NN)	0.10460000	0.06525657	0.04222260	0.03308495	0.02609508
GL-R1 (hyb.)	0.10459064	0.06450104	0.04141371	0.03162426	0.02770266
GL-R2 (hyb.)	0.10459064	0.06450091	0.04141371	0.03169532	0.02605176

Table 2: Energy levels $E(u)$ of computed energy minimizers for the reduced GL model of Section 5.1. First block: FE iterative solver with heuristic initial values φ_j . Second block: pure NN models GL-R1 (NN) and GL-R2 (NN). Third block: hybrid approach GL-R1 (hyb.) and GL-R2 (hyb.) with initial values from the NN.

The first block in Table 2 shows the energy levels $E(u)$ of the five computed minimizers using the classical FE iterative solver with the empirically chosen initial values φ_j , $j = 1, 2, 3, 4, 5$. The bold entries indicate, for each value of κ , the lowest energy achieved by this procedure; these serve as reference approximations for comparison with the NN and the hybrid approach.

For the NN-based method, we trained two models, GL-R1 and GL-R2. Their architectures and training strategies are summarized in Table 1. For both models, we use the SwiGLU block type and the main difference lies in the chosen learning rates and batch size. The training required approximately 3.5 hours for GL-R1 and 6.5 hours for GL-R2 on the HoreKa machine at KIT. The resulting energy levels of the pure NN approach are shown in the second block of Table 2. For $\kappa = 10, 25, 50$, the NN approximations yield energy levels close to, but slightly above, the reference values. In contrast, for $\kappa = 75, 100$ both models achieve lower energy levels than those obtained with the classical approach using the prescribed heuristic initial values. This suggests that the NN approach identifies better candidates for the global GL minimizer.

We then use our hybrid scheme, which takes the output of the NN as an initial guess for the classical FE iterative solver. The final energy levels after an additional iteration with the FE solver are shown in the third block of Table 2. Both models show comparable performance and consistently match or even outperform the reference approximations due to lower energy levels. In particular, they yield strictly lower energy states in several cases, indicating the successful identification of low energy states, which is the primary objective of our approach.

In Figure 2, we present the densities $|u|^2$ of the computed minimizers. In all cases, we observe the formation of a vortex pattern, namely the Abrikosov vortex lattice [1]. As κ is increased, the pattern becomes more intricate: the number of vortices grows while their size decreases. The minimizers computed with the classical FE approach and with the hybrid approach exhibit symmetric vortex patterns, whereas the pure




Fig. 2: Densities $|u|^2$ of computed minimizers for $\kappa = 10, 25, 50, 75, 100$ (left to right), corresponding to Table 2. Dark indicates values close to 1; light indicates values close to 0. First row: reference minimizer computed with the classical FE iterative solver and best heuristic initial values φ_j (best). Second and third row: pure NN models GL-R1 (NN) and GL-R2 (NN). Fourth and fifth row: hybrid approach GL-R1 (hyb.) and GL-R2 (hyb.) using initial values from the NN.

NN results show a certain displacement of the vortices within the lattice. At the same time, the pure NN approach still captures the overall structure of the pattern, as well as the approximate size and number of vortices. Moreover, for the same value of κ (e.g., $\kappa = 100$), even small differences in the computed energy correspond to different vortex configurations and therefore to different local minimizers. This suggests that, with the classical approach and heuristic initial values, the global GL minimizer is not recovered, while the hybrid approach appears to provide at least a strong candidate for the global minimizer.

Table 3 shows the iterations that the FE solver needed to converge up to the prescribed tolerance of 10^{-12} for both, the classical FE and the hybrid approach. As the table shows there is no clear indication whether either approach is favorable in that regard. Note that the numbers should only be compared for cases where the same minimizing state is achieved by both methods. While the performance is not the focus of our approach, we nonetheless implemented a refinement strategy to improve on the number of iterations needed after the NN stage. For this, we treat the trained NN as a pre-trained model that is then refined for a limited time and for the given value of κ . This can be interpreted as specializing the general model for this specific κ . In our tests, we observed clear indications that this procedure can indeed reduce the required number of iterations. However, a further detailed investigation of possibilities to reduce the number of iterations is beyond the scope of this work and

κ	10	25	50	75	100
φ_1	40	60	172	10279	10882
φ_2	199	339	1137	4373	24572
φ_3	66	1166	4034	8147	11022
φ_4	91	1214	4090	21718	17831
φ_5	86	379	21874	47145	37079
GL-R1 (hyb.)	52	1314	7483	11912	4537
GL-R2 (hyb.)	47	5541	7961	12452	24568

Table 3: Number of iterations to compute the minimizers with initial value $u^0 = \varphi_j$ and the hybrid-approach up to tolerance $\text{tol} = 10^{-12}$.

left for future research.

5.2. Neural networks for the full GL model. Next, we consider the full GL model (1.1) with the external magnetic field

$$(5.1) \quad \mathbf{h}_{\text{ext}}(x) = 2\sqrt{2}\pi \sin(\pi x_1) \sin(\pi x_2), \quad x = (x_1, x_2) \in \Omega,$$

and again choose the GL parameters as $\kappa = 10, 25, 50, 75, 100$. As before, we note that the NNs are not restricted to these values of κ , but instead are trained for a whole range.

κ	10	25	50	75	100
φ_1	0.10440628	0.08514562	0.04148069	0.04560323	0.02615202
φ_2	0.11500423	0.07371087	0.04185499	0.03776825	0.02854666
φ_3	0.10440628	0.06491247	0.04148069	0.03779534	0.02659625
φ_4	0.10440628	0.06491247	0.04148069	0.03446474	0.02660840
φ_5	0.10440628	0.08287191	0.04781070	0.03597329	0.02707063
GL-F1 (NN)	0.10461688	0.06576285	0.04211054	0.03220804	0.02639166
GL-F2 (NN)	0.10446705	0.06545582	0.04194608	0.03210672	0.02640819
GL-F1 (hyb.)	0.10440628	0.06445608	0.04148069	0.03171263	0.02615239
GL-F2 (hyb.)	0.10440628	0.06445609	0.04148069	0.03171263	0.02614767

Table 4: Energy levels $E(u, \mathbf{A})$ of computed energy minimizers for the full GL model of Section 5.1. First block: classical FE iterative solver with heuristic initial values φ_j for u and 0 for \mathbf{A} . Second block: pure NN models GL-F1 (NN) and GL-F2 (NN). Third block: hybrid approach GL-F1 (hyb.) and GL-F2 (hyb.) using initial values from the NN.

Similar to the previous experiment, the first block of Table 4 shows the computed minimizers obtained with the classical finite element iterative solver using heuristic initial values φ_j . The reference minimizers with the lowest energy levels are again highlighted in bold.

For the full model, we trained two NNs, denoted by GL-F1 and GL-F2, and identified architectures and training strategies that yield to minimizers of low energy. The corresponding parameters are listed in Table 1, and the resulting energy levels of the




Fig. 3: Densities $|u|^2$ of computed minimizers for $\kappa = 10, 25, 50, 75, 100$ (left to right), corresponding to Table 4. Dark indicates values close to 1; light indicates values close to 0. First row: reference minimizer computed with the classical FE iterative solver and best heuristic initial values φ_j (best). Second and third row: pure NN models GL-F1 (NN) and GL-F2 (NN). Fourth and fifth row: hybrid approach GL-F1 (hyb.) and GL-F2 (hyb.) using initial values from the NN.

computed GL minimizers are reported in the second block of Table 4. The training times of GL-F1 and GL-F2 were 7 and 4.5 hours, respectively, on the HoreKa machine at KIT. As in the reduced-model setting, the pure NN approaches produce energy levels close to, but usually slightly above, those of the reference minimizers. The only exception is the case $\kappa = 75$, where the networks find a minimizer with significantly lower energy, suggesting already a candidate for the global GL minimizer.

We then apply the hybrid approach to the full problem to refine the NN predictions. The corresponding energy levels for the hybrid approach using GL-F1 and GL-F2 are shown in the third block of Table 4. For $\kappa = 10$ and 50, the hybrid approach recovers the same energy states as the best classical methods, while for $\kappa = 25, 75$, and 100 it yields states with significantly lower energy. Hence, using NN predictions as initial values can improve the iterative solver and lead to lower-energy states.

Figure 3 shows plots of the densities $|u|^2$ of the computed minimizers. The observations are very similar to the case of the reduced model. The NNs yield rough approximations of the vortex pattern in the minimizers that, however, need to be refined by the iterative solver, i.e., the hybrid approach. In particular, considering the cases $\kappa = 25, 75, 100$ the hybrid approach compute a significant different vortex pattern of lower energy.

5.3. Fast-training neural networks. Finally, we consider a scenario in which, for a given GL model, the goal is to train a NN on demand, with a limited training time of about 1.5 hours. As a basis for the NNs, we use the same architecture and hyperparameters as for GL-F1 and train five different NNs by varying the initial seeds in the training process. We then exploit the hybrid approach to the given GL test model. This experiment evaluates the reliability and quality of the approximations achievable with the GL-F1 architecture and hyperparameters. Note that we do not use the trained model GL-F1, but instead train the same architecture from scratch with only 4 epochs. The obtained networks are termed GL-U*, $*$ $\in \{1, 2, 3, 4, 5\}$.

κ	10	25	50	75	100
φ_j (best)	0.10440628	0.06491247	0.04148069	0.03446474	0.02615202
GL-U1 (hyb.)	0.10440628	0.06445609	0.04148069	0.03176793	0.02615239
GL-U2 (hyb.)	0.10440628	0.06445608	0.04160989	0.03171263	0.02628949
GL-U3 (hyb.)	0.10440628	0.06445617	0.04160988	0.03171263	0.02614767
GL-U4 (hyb.)	0.10440628	0.06445608	0.04148069	0.03171263	0.02628949
GL-U5 (hyb.)	0.10440628	0.06445615	0.04148069	0.03171263	0.02660915

Table 5: Energy levels $E(u, \mathbf{A})$ of computed energy minimizers for the full GL model on the unit square domain. First block: classical FE iterative solver with heuristic initial values φ_j for u and 0 for \mathbf{A} . Second block: hybrid approach using the fast-trained GL-U* (hyb.), $*$ $\in \{1, 2, 3, 4, 5\}$.




Fig. 4: Densities $|u|^2$ of computed minimizers for $\kappa = 10, 25, 50, 75, 100$ (left to right), corresponding to Table 5. Dark indicates values close to 1; light indicates values close to 0. First row: reference minimizer computed with the classical FE iterative solver and best heuristic initial values φ_j (best). Second row: best hybrid approach using the fast-trained GL-U* (hyb.), $*$ $\in \{1, 2, 3, 4, 5\}$.

For the first test problem, we consider again the unit square and the external magnetic field \mathbf{h}_{ext} from (5.1). The energy levels of the resulting GL minimizers computed with the hybrid approach with the five different NNs are shown in Table 5. Again, the minimizers of the lowest energy are highlighted in bold. If there are multiple minimizers with the same energy, we highlight the one that requires the fewest iterations. We observe that there is only a minor variation in the energy levels using the hybrid method with GL-U*, trained with different seeds. Compared to the classical method with the heuristic initial values, the obtained energy levels are in favor of the NNs (see, e.g., the lower energy levels for $\kappa = 25$ or $\kappa = 75$). This can be

recognized also by considering the densities of the obtained minimizers that are shown in Figure 4. The different vortex patterns confirm that indeed a different minimizing state is obtained by the hybrid method.

κ	10	25	50	75	100
φ_j (best)	0.06958856	0.04353387	0.02864807	0.02261634	0.01901139
GL-L1 (hyb.)	0.06958856	0.04310817	0.02856779	0.02204774	0.01830038
GL-L2 (hyb.)	0.06958856	0.04310817	0.02854874	0.02204774	0.01830132
GL-L3 (hyb.)	0.06958856	0.04310817	0.02862300	0.02204774	0.01830037
GL-L4 (hyb.)	0.06958856	0.04310817	0.02862300	0.02204485	0.01829811
GL-L5 (hyb.)	0.06958856	0.04310817	0.02854874	0.02204774	0.01830038

Table 6: Energy levels $E(u, \mathbf{A})$ of computed energy minimizers for the full GL model on the L-shaped domain. First block: classical FE iterative solver with heuristic initial values φ_j for u and 0 for \mathbf{A} . Second block: hybrid approach using the fast-trained GL-L* (hyb.), $* \in \{1, 2, 3, 4, 5\}$.




Fig. 5: Densities $|u|^2$ of computed minimizers for $\kappa = 10, 25, 50, 75, 100$ (left to right), corresponding to Table 6. Dark indicates values close to 1; light indicates values close to 0. First row: reference minimizer computed with the classical FE iterative solver and best heuristic initial values φ_j (best). Second row: best hybrid approach using the fast-trained GL-L* (hyb.), $* \in \{1, 2, 3, 4, 5\}$.

For the second test problem, we replace the computational domain by an L-shaped domain, while keeping the external magnetic field \mathbf{h}_{ext} from (5.1). We employ the architecture and hyperparameters of GL-F1, but as before train the model from scratch, only adjusting the learning rate and its scheduling. We restrict the modifications of the model to these hyperparameters, since they are both essential for successful training and can be identified through a small number of test runs. In practice, tuning the learning rate and its schedule is typically necessary for successful and effective NN training. The obtained networks are termed GL-L*, $* \in \{1, 2, 3, 4, 5\}$ and its architecture and hyperparameters can be found in Table 1.

The energy levels of the computed minimizers by the hybrid approach for the five different models are shown in Table 6. For comparison, we again employ the classical FE iterative solver with the heuristic initial values φ_j and include the best obtained energy level in Table 6 as well. Again, we conclude that for all cases except for $\kappa = 10$ the NN-enhanced approach yields a minimizing state of lower energy than what is possible with the classical FE iterative solver and heuristic initial values. For

$\kappa = 10$ the same state is obtained. The difference in the energy states computed by the classical FE method and the hybrid method is also confirmed by comparing the densities, $|u|^2$, as shown in Figure 5.

Summarizing the two experiments, we conclude that the results achieved with the hybrid-approach are robust with respect to the random seeds used in the NN stage. Furthermore, the **GL-F1** seems to be reliable within the limitation of the experiments conducted. Therefore, the architecture of **GL-F1** can be considered as a good candidate for the NN stage in the hybrid approach when applying the method to other GL problems.

6. Conclusion. In this work, we proposed a neural network approach as well as a neural network-enhanced (hybrid) finite element method to compute minimizers of the Ginzburg–Landau energy (1.1). One of the main challenges in finding global minimizers is to choose good initial values for the minimization process as the results are highly sensitivity with respect to this choice. We have shown that neural networks can be used to approximate Ginzburg–Landau minimizers, often leading to already satisfactory energy levels. However, using these results as initial values or predictors for the finite element solver we can achieve even better results. Leveraging the computational advantage of neural networks exploiting massive parallelization on GPUs, we have further demonstrated that even for a short offline phase, compared to the computational time of the finite element solver, our neural network-enhanced approach can lead to minimizers of lower energy compared to a pure finite element approach with empirically chosen initial values. The proposed hybrid method does not necessarily imply faster convergence in terms of iteration steps, but we are able to find minimizers of lower energy, which we were not able to find with the pure finite element approach.

While the current neural network architecture achieves promising results and provides a good initial validation of the approach, we recognize that additional refinements and optimizations could further enhance its performance. However, these are beyond the scope of this study.

Acknowledgments. B. Dörich and R. Maier acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 258734477 – SFB 1173. C. Döding acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2047/1 – 390685813. M. Crocoll acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 545165789. Parts of this work were conducted during the authors’ stay at the Hausdorff Research Institute for Mathematics funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2047/2 – 390685813.

Further, we would like to thank Tim Buchholz for his valuable input on the MPI implementation of the FE solver and the research group IANM3 at KIT for including us in their computing project.

The authors gratefully acknowledge the computing time provided on the high-performance computer HoreKa by the National High-Performance Computing Center at KIT (NHR@KIT). This center is jointly supported by the Federal Ministry of Education and Research and the Ministry of Science, Research and the Arts of Baden-Württemberg, as part of the National High-Performance Computing (NHR) joint funding program (<https://www.nhr-verein.de/en/our-partners>). HoreKa is partly funded by the German Research Foundation (DFG).

This work is supported by the Helmholtz Association Initiative and Networking Fund on the HAICORE@KIT partition.

REFERENCES

- [1] A. A. ABRIKOSOV, *Nobel lecture: Type-II superconductors and the vortex lattice*, Rev. Mod. Phys., 76 (2004), pp. 975–979, <https://doi.org/10.1103/RevModPhys.76.975>.
- [2] J. ADLER AND O. ÖKTEM, *Solving ill-posed inverse problems using iterative deep neural networks*, Inverse Problems, 33 (2017), pp. 124007, 24, <https://doi.org/10.1088/1361-6420/aa9581>.
- [3] A. AFTALION, *On the minimizers of the Ginzburg–Landau energy for high kappa: the axially symmetric case*, Ann. Inst. H. Poincaré Anal. Non Linéaire, 16 (1999), pp. 747–772, [https://doi.org/10.1016/S0294-1449\(00\)88186-X](https://doi.org/10.1016/S0294-1449(00)88186-X).
- [4] J. ANSEL ET AL., *PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation*, in 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS ’24), ACM, Apr. 2024, <https://doi.org/10.1145/3620665.3640366>.
- [5] I. A. BARATTA ET AL., *DOLFINx: The next generation FEniCS problem solving environment*, Dec. 2023, <https://doi.org/10.5281/zenodo.10447666>.
- [6] M. BLUM, C. DÖDING, AND P. HENNING, *Vortex-capturing multiscale spaces for the Ginzburg–Landau equation*, Multiscale Model. Simul., 23 (2025), pp. 339–373, <https://doi.org/10.1137/24M1664538>.
- [7] T. CHAUMONT-FRELET AND P. HENNING, *The pollution effect for FEM approximations of the Ginzburg–Landau equation*, Math. Comp., (2026), p. 49pp., <https://doi.org/10.1090/mcom/4183>. Online first.
- [8] Y. N. DAUPHIN, A. FAN, M. AULI, AND D. GRANGIER, *Language Modeling with Gated Convolutional Networks*, 70 (2017), pp. 933–941, <https://proceedings.mlr.press/v70/dauphin17a.html>.
- [9] C. DÖDING, B. DÖRICH, AND P. HENNING, *A multiscale approach to the stationary Ginzburg–Landau equations of superconductivity*, CRC 1173 Preprint 2024/21, Karlsruhe Institute of Technology, sep 2024, <https://doi.org/10.5445/IR/1000174322>.
- [10] C. DÖDING AND P. HENNING, *The Ginzburg–Landau equations: Vortex states and numerical multiscale approximations*, ArXiv Preprint, 2511.19540 (2025).
- [11] B. DÖRICH, *Approximation of minimizers of the Ginzburg–Landau energy in non-convex domains*, CRC 1173 Preprint 2025/44, Karlsruhe Institute of Technology, 2025, <https://doi.org/10.5445/IR/1000183843>.
- [12] B. DÖRICH AND P. HENNING, *Error bounds for discrete minimizers of the Ginzburg–Landau energy in the high- κ regime*, SIAM J. Numer. Anal., 62 (2024), pp. 1313–1343, <https://doi.org/10.1137/23M1560938>.
- [13] Q. DU, M. D. GUNZBURGER, AND J. S. PETERSON, *Analysis and approximation of the Ginzburg–Landau model of superconductivity*, SIAM Rev., 34 (1992), pp. 54–81, <https://doi.org/10.1137/1034003>.
- [14] Q. DU, M. D. GUNZBURGER, AND J. S. PETERSON, *Modeling and analysis of a periodic Ginzburg–Landau model for type-II superconductors*, SIAM J. Appl. Math., 53 (1993), pp. 689–717, <https://doi.org/10.1137/0153035>.
- [15] Q. DU AND L. JU, *Approximations of a Ginzburg–Landau model for superconducting hollow spheres based on spherical centroidal Voronoi tessellations*, Math. Comp., 74 (2005), pp. 1257–1280, <https://doi.org/10.1090/S0025-5718-04-01719-3>.
- [16] Q. DU, R. A. NICOLAIDES, AND X. WU, *Analysis and convergence of a covolume approximation of the Ginzburg–Landau model of superconductivity*, SIAM J. Numer. Anal., 35 (1998), pp. 1049–1072, <https://doi.org/10.1137/S0036142996302852>.
- [17] W. E AND B. YU, *The Deep Ritz Method: A Deep Learning-Based Numerical Algorithm for Solving Variational Problems*, Commun. Math. Stat., 6 (2018), pp. 1–12.
- [18] I. GOODFELLOW, Y. BENGIO, AND A. COURVILLE, *Deep Learning*, MIT Press, 2016.
- [19] P. GOYAL, P. DOLLÁR, R. GIRSHICK, P. NOORDHUIS, L. WESOLOWSKI, A. KYROLA, A. TULLOCH, Y. JIA, AND K. HE, *Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour*, 2018, <https://arxiv.org/abs/1706.02677>.
- [20] K. JORDAN, Y. JIN, V. BOZA, J. YOU, F. CESISTA, L. NEWHOUSE, AND J. BERNSTEIN, *Muon: An optimizer for hidden layers in neural networks*, 2024, <https://kellerjordan.github.io/posts/muon/>.
- [21] D. P. KINGMA AND J. BA, *Adam: A Method for Stochastic Optimization*, in International

- Conference on Learning Representations, 2014.
- [22] J. LIU ET AL., *Muon is scalable for LLM Training*, ArXiv Preprint, 2502.16982 (2025), <https://arxiv.org/abs/2502.16982>.
- [23] H. K. ONNES, *Further experiments with liquid helium. c. on the change of electric resistance of pure metals at very low temperatures, etc. iv. the resistance of pure mercury at helium temperatures*, Comm. Phys. Lab. Univ. Leiden, 120 (1911).
- [24] D. PETERSEIM, J.-F. PIETSCHMANN, J. PÜSCHEL, AND K. RUESS, *Neural network acceleration of iterative methods for nonlinear Schrödinger eigenvalue problems*, J. Comput. Appl. Math., 485 (2026), p. Paper No. 117414, <https://doi.org/10.1016/j.cam.2026.117414>.
- [25] E. POLAK AND G. RIBIÈRE, *Note sur la convergence de méthodes de directions conjuguées*, Rev. Française Informat. Recherche Opérationnelle, 3 (1969), pp. 35–43.
- [26] M. RAISSI, P. PERDIKARIS, AND G. KARNIADAKIS, *Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations*, J. Comput. Phys., 378 (2019), pp. 686–707, <https://doi.org/10.1016/j.jcp.2018.10.045>.
- [27] E. SANDIER AND S. SERFATY, *Vortices in the magnetic Ginzburg-Landau model*, vol. 70 of Progress in Nonlinear Differential Equations and their Applications, Birkhäuser Boston, Inc., Boston, MA, 2007.
- [28] E. SANDIER AND S. SERFATY, *Vortices in the magnetic Ginzburg-Landau model*, vol. 70 of Progress in Nonlinear Differential Equations and their Applications, Birkhäuser Boston, Inc., Boston, MA, 2007.
- [29] E. SANDIER AND S. SERFATY, *From the Ginzburg-Landau model to vortex lattice problems*, Comm. Math. Phys., 313 (2012), pp. 635–743, <https://doi.org/10.1007/s00220-012-1508-x>.
- [30] S. SERFATY, *Stable configurations in superconductivity: uniqueness, multiplicity, and vortex-nucleation*, Arch. Ration. Mech. Anal., 149 (1999), pp. 329–365, <https://doi.org/10.1007/s002050050177>.
- [31] S. SERFATY AND E. SANDIER, *Vortex patterns in Ginzburg-Landau minimizers*, in XVIth International Congress on Mathematical Physics, World Sci. Publ., Hackensack, NJ, 2010, pp. 246–264, https://doi.org/10.1142/9789814304634_0014.
- [32] N. SHAZEER, *GLU variants improve transformer*, ArXiv Preprint, 2002.05202 (2020).
- [33] H. TOUVRON, M. CORD, A. SABLAYROLLES, G. SYNNAEVE, AND H. JEGOU, *Going deeper with Image Transformers*, in 2021 IEEE/CVF International Conference on Computer Vision (ICCV), Los Alamitos, CA, USA, Oct. 2021, IEEE Computer Society, pp. 32–42, <https://doi.org/10.1109/ICCV48922.2021.00010>.
- [34] S. VENKATARAMAN AND B. AMOS, *Neural Fixed-Point Acceleration for Convex Optimization*, in 8th ICML Workshop on Automated Machine Learning (AutoML), 2021, <https://openreview.net/forum?id=Vxbpb6XvGgH>.
- [35] H. ZHANG, Y. N. DAUPHIN, AND T. MA, *Fixup initialization: Residual learning without normalization*, ArXiv Preprint, 1901.09321 (2019), <https://arxiv.org/abs/1901.09321>.