

## Mathematical Analysis of Neural Networks and Generative AI

**D. William John Victor**  
Associate Professor,  
Dept. of Mathematics,  
AMAL College, Anakapalli

**Janga. Sambraiyam**  
Lecturer,  
Dept. of Mathematics,  
AMAL College, Anakapalli.

**Abstract:** This overview article highlights the critical role of mathematics in artificial intelligence (AI), emphasizing that mathematics provides tools to better understand and enhance AI systems. Conversely, AI raises new problems and drives the development of new mathematics at the intersection of various fields. This article focuses on the application of analytical and probabilistic tools to model neural network architectures and better understand their optimization. Statistical questions (particularly the generalization capacity of these networks) are intentionally set aside, though they are of crucial importance. We also shed light on the evolution of ideas that have enabled significant advances in AI through architectures tailored to specific tasks, each echoing distinct mathematical techniques. The goal is to encourage more mathematicians to take an interest in and contribute to this exciting field

**Key words:** Minimization, data, generative, function, Analysis.

### Supervised Learning

Recent advancements in artificial intelligence have mainly stemmed from the development of neural networks, particularly deep networks. The first significant successes, after 2010, came from supervised learning, where training pairs  $(x^i, y^i)_i$  are provided, with  $x^i$  representing data (e.g., images or text) and  $y^i$

corresponding labels (typically, classes describing the content of the data). More recently, spectacular progress has been achieved in unsupervised learning, where labels  $y_i$  are unavailable, thanks to techniques known as generative AI. These methods will be discussed in Section 4.

**Empirical Risk Minimization.** In supervised learning, the objective is to construct a function  $f_\theta(x)$ , dependent on parameters  $\theta$ , such that it approximates the data well:  $y^i \approx f_\theta(x^i)$ . The dominant paradigm involves finding the parameters  $\theta$  by minimizing an empirical risk function, defined as:

$$\theta \min E(\theta) := \sum_i \ell(f_\theta(x^i), y^i), \quad (1)$$

where  $\ell$  is a loss function, typically  $\ell(y, y') = \|y - y'\|^2$  for vector-valued  $(y^i)_i$ . Optimization of  $E(\theta)$  is performed using gradient descent:



$$\theta_{t+1} = \theta_t - \tau \nabla E(\theta_t), \quad (2)$$

where  $\tau$  is the step size. In practice, a variant known as stochastic gradient descent [23] is used to handle large datasets by randomly sampling a subset at each step  $t$ . A comprehensive theory for the convergence of this method exists in cases where  $f_\theta$  depends linearly on  $\theta$ , as  $E(\theta)$  is convex. However, in the case of deep neural networks,  $E(\theta)$  is non-convex, making its theoretical analysis challenging and still largely unresolved. In some instances, partial mathematical analyses provide insights into the observed practical successes and guide necessary modifications to improve

convergence. For a detailed overview of existing results, refer to [2].

**Automatic Differentiation.**

Computing the gradient  $E(\theta)$  is fundamental. Finite difference methods or traditional differentiation of composite functions would be too costly, with complexity on the order of  $O(DT)$ , where  $D$  is the dimensionality of  $\theta$  and  $T$  is the computation time of  $f_\theta$ . The advances in deep network optimization rely on the use of automatic differentiation in reverse mode [13], often referred to as “backpropagation of the gradient,” with complexity on the order of  $O(T)$ .

**Two-Layer Neural Networks Multi-Layer Perceptrons.** A multi-layer network, or multi-layer perceptron (MLP) [25], computes the function  $f_\theta(x) = x_L$  in  $L$  steps (or layers) recursively starting from  $x_0 = x$ :

$$x_{\ell+1} = \sigma(W_\ell x_\ell + b_\ell), \quad (3)$$

where  $W_\ell$  is a weight matrix,  $b_\ell$  a bias vector, and  $\sigma$  a non-linear activation function,

such as the sigmoid  $\sigma(s) = \frac{e^s}{1+e^s}$ , which is bounded, or ReLU (Rectified Linear Unit),

$\sigma(s) = \max(s, 0)$ , which is unbounded. The parameters to optimize are the weights and biases  $\theta = (W_\ell, b_\ell)_\ell$ . If  $\sigma$  were linear, then  $f_\theta$  would remain linear regardless of the number

of layers. The non-linearity of  $\sigma$  enriches the set of functions representable by  $f_\theta$ , by increasing both the width (dimension of intermediate vectors  $x_k$ ) and the depth (number  $L$  of layers).

**Two-Layer Perceptrons and Universality.** The mathematically best-understood case is that of  $L = 2$  layers. Denoting  $W_1 = (v_k)^n$  and  $W_2^\top = (u_k)^n$  as the  $n$  rows

and columns of the two matrices (where  $n$  is the network width), we can write

$f_\theta$  as a sum of contributions from its  $n$  neurons:

$$f_\theta(x) = \frac{1}{n} \sum_{k=1}^n u_k \sigma(\langle v_k, x \rangle + b_k). \quad (4)$$

The parameters are  $\theta = (u_k, v_k, b_k)_{k=1}^n$ . Here, we added a normalization factor  $1/n$  (to

later study the case  $n \rightarrow +\infty$ ) and ignored the non-linearity of the second layer. A classical result by Cybenko [8] shows that these functions  $f_\theta$  can uniformly approximate any continuous function on a compact domain. This result is similar to the Weierstrass theorem for polynomial approximation, except that (4) defines, for a fixed  $n$ , a non-linear function space. The elegant proof by Hornik [15] relies on the Stone-Weierstrass theorem, which implies the result for  $\sigma = \cos$  (since (4) defines an algebra of functions when  $n$  is arbitrary). The proof is then completed by uniformly approximating a cosine in 1-D using a sum of sigmoids or piecewise affine functions (e.g., for  $\sigma = \text{ReLU}$ ).

**Mean-Field Representation.** This result is, however, disappointing, as it does not specify how many neurons are needed to achieve a given approximation error. This is impossible without adding assumptions about the regularity of the function to approximate. Barron's fundamental result [4] introduces a Banach space of regular functions, defined by the semi-norm  $\|f\|_B := \int \hat{f}(\omega) \omega d\omega$ , where  $\hat{f}$  denotes the Fourier transform of  $f$ . Barron shows that, in the  $L^2$  norm, for any  $n$ , there exists a network  $f_\theta$  with  $n$  neurons such that the approximation error on a compact  $\Omega$  of radius  $R$  is of the order

$\|f - f_\theta\|_{L^2(\Omega)} = O\left(\frac{R \|f\|_B}{n}\right)$ . This result is remarkable because it avoids the "curse of

dimensionality": unlike polynomial approximation, the error does not grow exponentially with the dimension  $d$  (although the dimension affects the constant  $\|f\|_B$ ).

Barron's proof relies on a "mean-field" generalization of (4), where a distribution (a probability measure)  $\rho$  is considered over the parameter space  $(u, v, b)$ , and the network

$$F_\rho(x) := \int u \sigma(\langle v, x \rangle + b) d\rho(u, v, b). \quad (5)$$

A finite-size network (4),  $f_\theta = F_\rho$ , is obtained with a discrete measure

$$\hat{\rho} = \frac{1}{n} \sum_{k=1}^n \delta_{(u_k, v_k, b_k)}.$$



An advantage of the representation (5) is that  $F_\rho$  depends linearly on  $\gamma$ , which is crucial for Barron's proof. This proof uses a probabilistic Monte Carlo-like method (where the error decreases as  $1/n$  as desired): it involves constructing a distribution  $\rho$  from  $f$ , then sampling a discrete measure  $\hat{\rho}$  whose parameters  $(u_k, v_k, b_k)_k$  are distributed according to  $\rho$ .

**Wasserstein Gradient Flow.**

In general, analyzing the convergence of the optimization (2) is challenging because the function  $E$  is non-convex. Recent analyses

have shown that when the number of neurons  $n$  is large, the dynamics are not trapped in a local minimum. The fundamental result by Chizat and Bach [7] is based on the fact that the distribution  $\rho_t$  defined by gradient descent (2), as  $\tau_t \rightarrow 0$ , follows a gradient flow in the space of probability distributions equipped with the optimal transport distance. These Wasserstein gradient flows, introduced by [16] and studied extensively in the book [1], satisfy a McKean-Vlasov-type partial differential equation

$$\partial_t \rho_t + \text{div}(\rho_t V(\rho)) = 0, \tag{6}$$

where  $V(\rho)(x)$  is a vector field depending on  $\rho$ , which can be computed explicitly from the data  $(x^i, y^i)_i$ . The result by Chizat and Bach can be viewed both as a PDE result (convergence to an equilibrium of a class of PDEs with a specific vector field  $V(\rho)$ ) and as a machine learning result (successful training of a two-layer network via gradient descent when the number of neurons is sufficiently large).

**Very Deep Networks**

The unprecedented recent success of neural networks began with the work of [17], which demonstrated that deep networks, when trained on large datasets, achieve unmatched

performance. The first key point is that to achieve these performances, it is necessary to use weight matrices  $W_k$  that exploit the structure of the data. For images, this means using convolutions [18]. However, this approach is not sufficient for extremely deep networks, with the number of layers  $L$  reaching into the hundreds

**Residual Networks.** The major breakthrough that empirically demonstrated that network performance increases with  $L$  was the introduction of residual connections, known as ResNet [14]. The main idea is to ensure that, for most layers  $l$ , the dimensions of  $x^l$  and  $x^{l+1}$  are identical, and to replace (3) with  $L$  steps



$$x_{\ell+1} = x_{\ell} + U_{\ell} \tau \sigma(V_{\ell} x_{\ell} + b_{\ell}), \quad (7)$$

where  $U_{\ell}, V_{\ell} \in \mathbb{R}^{n \times d}$  are weight matrices, with  $n$  being the number of neurons per layer (which, as in (4), can be increased to enlarge the function space). The intuition for the success of (7) is that this formula allows, unlike (3), for steps that are small deformations near identity mappings. This makes the network  $f_{\theta}(x_0) = x_L$  obtained after  $L$  steps well-posed even when  $L$  is

large, and it can be rigorously proven [20] that this well-posedness is preserved during optimization via gradient descent (2).

Neural Differential Equation. As  $L$  approaches  $+\infty$ , (7) can be interpreted as a discretization of an ordinary differential equation

$$\frac{dx_s}{ds} = U_s \tau \sigma(V_s x_s + b_s), \quad (8)$$

where  $s \in [0, 1]$  indexes the network depth. The network  $f_{\theta}(x_0) := x_1$  maps the initialization  $x_{s=0}$  to the solution  $x_{s=1}$  of (8) at time  $s = 1$ . The parameters are  $\theta = (U_s, V_s, b_s)_{s \in [0,1]}$ . This formalization, referred to as a *neural ODE*, was initially introduced in [6] to leverage tools from adjoint equation theory to reduce the memory cost of computing the gradient  $\nabla E(\theta)$  during backpropagation. It also establishes a connection to control theory, as training via gradient descent (2) computes an *optimal control*  $\theta$  that interpolates between the data  $x^i$  and labels  $y^i$ . However, the specificity of learning theory compared to control theory lies in the goal of computing such control using gradient descent (2). To date, no detailed results exist on this. Nonetheless, in his thesis, Raphael Barboni [3] demonstrated that if the network is initialized near an

interpolating network, gradient descent converges to it

**Generative AI for Vector Data**

**Self-Supervised Pre-Training.** The remarkable success of large generative AI models for vector data, such as images (often referred to as “diffusion models”), and for text (large language models or LLMs), relies on the use of new large-scale network architectures and the creation of new training tasks known as “self-supervised”. Manually defining the labels  $y^i$  through human intervention is too costly, so these are calculated automatically by solving simple tasks. For images, these involve denoising tasks, while for text, they involve predicting the next word. A key advantage of these simple tasks (known as pre-training tasks) is that it becomes possible to use a pre-trained network  $f_{\theta}$  generatively. Starting from an image composed of pure noise and iterating on the



network, one can randomly generate a realistic image [27]. Similarly, for text, starting from a prompt and sequentially predicting the next word, it is possible to generate text, for example, to answer questions [5]. We will first describe the case of vector data generation, such as images, and address LLMs for text in the following section

**Sampling as an Optimization Problem.**

Generative AI models aim to generate (or “sample”) random vectors  $x$  according to a distribution  $\beta$ , which is learned from a large training dataset  $(x^i)_i$ . This distribution is obtained by “pushing forward” a

This problem is challenging to optimize because computing  $D$  is expensive, and  $\beta$  must be approximated from the data  $(x^i)_i$ .

**Flow-Based Generation.** Recent successes, particularly

where  $g_\theta : (x, s) \in \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ . The input space includes an additional temporal dimension  $s$ , compared to  $f_\theta$ . The most effective neural networks for images are notably U-Nets [24]. The central mathematical

reference distribution  $\alpha$  (most commonly an isotropic Gaussian distribution  $\alpha = N(0, Id)$ ) through a neural network  $f_\theta$ . Specifically, if  $X \sim \alpha$  is distributed according to  $\alpha$ , then  $f_\theta(X) \sim \beta$  follows the law  $\beta$ , which is denoted as  $(f_\theta)_\# \alpha = \beta$ , where  $\#$  represents the pushforward operator (also known as the image measure in probability theory).

Early approaches to generative AI, such as “Generative Adversarial Networks” (GANs) [12], attempted to directly optimize a distance  $D$  between probability distributions (e.g., an optimal transport distance [22]):

$$\min_{\theta} D((f_\theta)_\# \alpha, \beta). \tag{9}$$

in the generation of vector data, involve neural networks  $f_\theta$  that are computed iteratively by integrating a flow [21], similar to a neural differential equation (8):

$$\frac{df(x)}{ds} = g_\theta(x, s)$$

$$\text{where } \frac{dx}{ds} = g_\theta(x, s), \tag{10}$$

question is how to replace (9) with a simpler optimization problem when  $f_\theta$  is computed by integrating a flow (10). An extremely elegant solution was first proposed in the context of diffusion models [27] (corresponding to the specific case  $\alpha = N(0, Id)$ ) and later generalized



under the name “flow matching” [19].

The main idea is that if  $x_0 \sim \alpha$  is initialized randomly, then  $x_s$ , the solution at time  $s$  of (10), follows a distribution

$$\partial_s \alpha_s + \text{div}(\alpha_s v_s) = 0,$$

where the vector field is defined as  $v_s(x) := g_\theta(x, s)$ . This equation is similar to the evolution of neuron distributions during optimization (6), but it is simpler because  $v_s$  does not depend on the distribution  $\alpha_s$ , making the equation linear in  $\alpha_s$ .

**Denoising Pre-Training.**

The question is how to find a vector field  $v_s(x) = g_\theta(x, s)$  such that  $\alpha_1 = \beta$ , i.e., the final distribution matches the desired one. If  $v_s$  is known, the distribution  $\alpha_s$  is uniquely determined. The key idea is to reason in reverse: starting from an interpolation  $\alpha_s$  satisfying  $\alpha_0$

$\alpha_s$  that interpolates between  $\alpha_0 = \alpha$  and  $\alpha_1$ , which is desired to match  $\beta$ . This distribution satisfies a conservation equation:

$$(11)$$

$= \alpha$  and  $\alpha_1 = \beta$ , how can we compute a vector field  $v_s$  such that (11) holds? There are infinitely many possible solutions (since  $v_s$  can be modified by a conservative field), but for certain specific interpolations  $\alpha_s$ , remarkably simple expressions can be derived.

One example is the interpolation obtained via barycentric averaging: take a pair  $x_0 \sim \alpha$  and  $x_1 \sim \beta$ , and define  $\alpha_s$  as the distribution of  $(1 - s)x_0 + sx_1$ . It can be shown [19] that an admissible vector field is given by a simple conditional expectation:

$$v_s(x) = E_{x_0 \sim \alpha, x_1 \sim \beta} x_1 - x_0 \mid (1 - s)x_0 + sx_1 = x. \quad (12)$$

The key advantage of this formula is that the conditional mean corresponds to a linear regression, which can be approximated using a neural network  $v_s \approx g_\theta(\cdot, s)$ . The

expectation over  $x_1 \sim \beta$  can be replaced by a sum over training data  $(x^i)_i$ , leading to the following optimization problem:



$$\min_{\theta} E(\theta) := \int_0^1 \mathbb{E}_{x_0} \alpha \left\| x^i - x^0 - g_{\theta}((1-s)x_0 + sx^i, s) \right\|^2 ds.$$

This function  $E(\theta)$  is an empirical risk function (1), and similar optimization techniques are used to find an optimal  $\theta$ . In particular, stochastic gradient descent efficiently handles the problem (12) to an unsupervised no labels  $y^i$ , but an artificial supervision task is created by adding random noise  $x_0$  to the data  $x^i$ . This task is called denoising. We will now see that, for textual data, a different pre-training task is used

### Generative AI for Text

Tokenization and Next-Word Prediction. Generative AI methods for text differ from those used for vector data generation. The neural network architectures are different (they involve transformers, as we will describe), and the pre-training method is based not on denoising but on next-word prediction [28]. It

is worth noting that these transformer neural networks are now also used for image generation [9], but specific aspects related to the causal nature of text remain crucial. The first preliminary step, called “tokenization”, consists of transforming the input text into a sequence of vectors  $X = (x[1], \dots, x[P])$ , where the number  $P$  is variable (and may increase, for instance, when generating text from a prompt). Each token  $x[p]$  is a vector that encodes a group of letters, generally at the level of a syllable. The neural network  $x = f_{\theta}(X)$  is applied to all the tokens and predicts a new token  $x$ . During training, a large corpus of text  $(X^i)_i$  is available. If we denote by  $\tilde{X}^i$  the text  $X^i$  with the last token  $x^i$  removed, we minimize an empirical risk for prediction:

$$\min_{\theta} E(\theta) := \sum_i \ell(f_{\theta}(\tilde{X}^i), x^i),$$

which is exactly similar to (1). When using a pre-trained network  $f_{\theta}$  for text generation, one starts with a prompt  $X$  and iteratively adds a new token in an auto-regressive manner:  $X \rightarrow [X, f_{\theta}(X)]$ .

### Transformers and Attention.

The large networks used for text generation tasks are Transformer networks [10]. Unlike ResNet (7), these networks  $f_{\theta}$  no longer operate on a single vector  $x$ , but on a set of vectors  $X = (x[1], \dots, x[P])$  of size  $P$ . In



Transformers, the ResNet layer (7), operating on a single vector, is replaced by an *attention* layer, where all

tokens interact through a barycentric combination of vectors  $(V x[q])_q$  where  $V$  is a matrix:

$$A_{\omega}(X)_p := \sum_q M_{p,q} (V x[q]), \quad \text{where } M_{p,q} := \frac{e^{\langle Qx[p], Kx[q] \rangle}}{\sum_{\ell} e^{\langle Qx[p], Kx[\ell] \rangle}}$$

with  $\omega := (Q, K, V)$ .

(13)

The coefficients  $M_{p,q}$  are computed by normalized correlation between  $x[p]$  and  $x[q]$ , depending on two parameter matrices  $Q$  and  $K$ . A Transformer  $f_{\theta}(X)$  is then defined similarly to ResNet (7), iterating  $L$  attention layers with residual connections:

$$X_{\ell+1} = X_{\ell} + \frac{1}{L} A_{\omega \ell}(X_{\ell}). \quad - (14)$$

The parameters  $\theta$  of the Transformer  $f_{\theta}(X_0) = X_L$  with  $L$  layers are  $\theta = (\omega_{\ell} = (Q_{\ell}, K_{\ell}, V_{\ell}))^{L-1}$ .

This description is simplified: in practice, a Transformer network also integrates normalization layers and MLPs operating independently on **Mean-Field Representation of Attention**. The attention mechanism (13) can be viewed as a system of interacting particles, and (14) describes the evolution of

each token. For text applications, attention must be causal, imposing  $M_{i,j} = 0$  for  $j > i$ . This constraint is essential for recursively generating text and ensuring the next-word prediction task is meaningful.

tokens across depth. As with ResNet (8), in the limit  $L \rightarrow +\infty$ , one can consider a system of coupled ordinary differential equations:

$$\frac{dX_s}{ds} = A_{\omega s} \cdot (X_s) \quad (15)$$



distribution  $\mu_s := \frac{1}{p} \sum_p \delta_{x_S}[p]$  over the token space. This perspective was adopted in Michael Saneer's thesis [20], which writes attention as,

$$A_\omega(x) := \int \frac{e^{\langle Qx, Kx' \rangle} V x'}{e^{\langle Qx, Kx' \rangle} d\mu(x')} \text{ where } \omega := (Q, K, V).$$

The particle system (15) then becomes a conservation equation for advection by the vector field  $A_{\omega_S}(\mu)$ :

$$\partial_s \mu_s + \text{div}(\mu_s A_{\omega_S}(\mu_s)) = 0. \quad (16)$$

Surprisingly, this yields a McKean-Vlasov-type equation, similar to the one describing the training of two-layer MLPs (6), but with the velocity field  $A_\omega(\mu)(x)$  replacing  $(\rho)(u, v, b)$ . However, here the evolution occurs in the token space  $x$  rather than in the neuron space  $(u, v, b)$ , and the evolution variable corresponds to depth  $s$ , not the

optimization time  $t$ .

Unlike (6), the evolution (16) is not a gradient flow in the Wasserstein metric [26]. Nonetheless, in certain cases, this evolution can be analyzed, and it can be shown that the measure  $\mu_s$  converges to a single Dirac mass [11] as  $s \rightarrow \infty$ : the tokens tend to cluster. Better understanding these evolutions, as well as the optimization of parameters  $\theta = (Q_s, K_s, V_s)_{s \in [0,1]}$  via gradient descent (2), remains an open problem. This problem can be viewed as a control problem for the PDE (16).



## Conclusion

Mathematics plays a critical role in understanding and improving the performance of deep architectures while presenting new theoretical challenges. The emergence of Transformers and generative AI raises immense mathematical problems, particularly for better understanding the training of these networks and exploring the structure of "optimal" networks. One essential question remains: whether an LLM merely interpolates training data or is capable of genuine reasoning. Moreover, issues of resource efficiency and privacy in AI system development demand significant theoretical advancements, where mathematics will play a pivotal role. Whether designing resource-efficient models, ensuring compliance with ethical standards, or exploring the fundamental limits of these systems, mathematics is poised to be an indispensable tool for the future of artificial intelligence

## References

- [1] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.
- [2] Francis Bach. *Learning theory from first principles*. MIT press, 2024.
- [3] Raphaël Barboni, Gabriel Peyré, and François-Xavier Vialard. Understanding the training of infinitely deep and wide resnets with conditional optimal transport. *arXiv preprint arXiv:2403.12887*, 2024.
- [4] Andrew R Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [6] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- [7] Lenaïc Chizat and Francis Bach. On the global convergence of gradient descent for overparameterized models using optimal transport. *Advances in neural information processing systems*, 31, 2018.
- [8] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- [9] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [10] A. Vaswani et al. Attention is



- all you need. *Advances in Neural Information Processing Systems*, 2017.
- [11] Borjan Geshkovski, Cyril Letrouit, Yury Polyanskiy, and Philippe Rigollet. The emergence of clusters in self-attention dynamics. *Advances in Neural Information Processing Systems*, 36, 2024.
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Advances in neural information processing systems 27*, 2014.
- [13] Andreas Griewank and Andrea Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, 2008.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [15] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- [16] Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the fokker-planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [18] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [19] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *Proc. ICLR 2023*, 2023.
- [20] Pierre Marion, Yu-Han Wu, Michael E Sander, and Gérard Biau. Implicit regularization of deep residual networks towards neural odes. *Proc. ICLR'23*, 2023.
- [21] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- [22] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [23] Herbert Robbins and Sutton



- Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [24] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- [25] Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [26] Michael E Sander, Pierre Ablin, Mathieu Blondel, and Gabriel Peyré. Sinkformers: Transformers with doubly stochastic attention. In *International Conference on Artificial Intelligence and Statistics*, pages 3515–3530. PMLR, 2022.
- [27] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.