# Modelling variation in the forward EMG model.

**Dimitrios Halatsis** [*]
Imperial College London
d.halatsis@imperial.ac.uk

**Alexander Clarke**
Imperial College London
a.clarke18@imperial.ac.uk

**Dario Farina**
Imperial College London
d.farina@imperial.ac.uk

## Abstract

Traditional finite element methods (FEM) face limitations when simulating dynamic, naturalistic movements due to the need for recalculations as muscle geometry changes. Our approach leverages PINNs to implicitly represent a continuous range of solutions for the volume conduction equation, parameterized by the pinnation angle of muscle fibers. We demonstrate that our method significantly reduces model size and computational time while maintaining high accuracy. The neural network model generalizes well across a range of pinnation angles, offering a promising solution for efficient and dynamic EMG simulations.

## 1 Introduction

Biophysical simulation plays a crucial role in contemporary biomedical research and engineering, offering a cost-effective means of conducting experiments and testing hypotheses prior to physical implementation (Gerstner et al., 2012). This is particularly vital in the field of electromyography (EMG) research, where anatomically precise simulations are essential for developing algorithms in data-scarce environments (Clarke & Farina, 2023; Merletti et al., 1999; Mamidanna & Farina).

Obtaining anatomically accurate simulations requires detailed modeling, which entails significant computational costs. These costs arise from the need to solve volume conduction partial differential equations (PDEs) within complex and densely structured domains that represent the human body's physical components. Finite element methods (FEM) are typically employed to solve these PDEs, as they can effectively handle intricate geometries (Maksymenko et al., 2023).

However, FEM faces limitations when simulating dynamic, naturalistic movements. FEM is highly reliable when the volume conductor remains static, but during dynamic scenarios, the geometry of the volume conductor changes as fibers contract and muscles move, necessitating computationally expensive recalculations of the solutions. Currently, the most effective approaches involve dividing movements into a series of static stages and solving them individually (Pereira Botelho et al., 2019).

With the advancement of data-driven methods such as Physics-Informed Neural Networks (PINNs), neural implicit representations and operator networks (Sitzmann et al., 2020; Sirignano & Spiliopoulos, 2018; Raissi et al., 2019; Lu et al., 2019), neural networks (NNs) have gained the ability to learn implicit representations of PDE solutions and even directly solve PDEs. This results in significantly faster simulations and the capacity to represent a continuous range of solutions, even when trained on a finite set of samples. Neural networks also amortize computational time during training, leading to extremely low inference times, making them an ideal fit for simulating dynamic movements. In the context of EMG simulation, Ma et al. (2024) utilized a neural network to generate motor unit

[*] Some information on the author https://dhalatsis.github.io/

D3S3: Data-driven and Differentiable Simulations, Surrogates, and Solvers @ NeurIPS 2024.

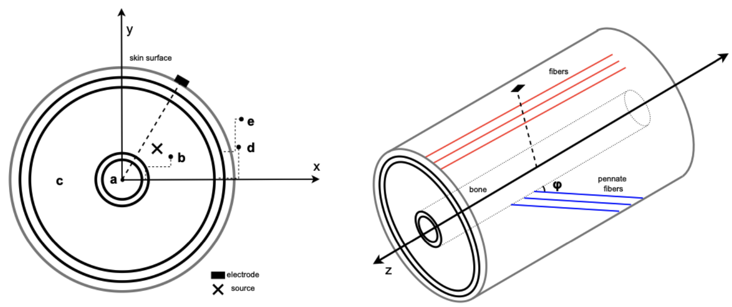


Figure 1: Cylindrical volume conductor model. Lef: Different sections correspond to different materials, *a* and *b* are cancelous and cortical bone, *c* is muscle, *d* and *e* are fat and skin respectively. Right: Muscle fibers are displayed with color, red parallel fibers and blue pennate fibers.

action potential waveforms across various simulation parameters, resulting in a much faster simulator capable of dynamically producing solutions for different limb positions.

In this work, we focus on creating a solver capable of representing solutions to the volume conduction equation for varying volume conductor geometries. We introduce a simple yet ubiquitous variation—the pinnation angle of muscle fibers (Farina et al., 2004a; Mesin & Farina, 2004)- and train a PINN to implicitly represent the function continuous range of solutions.

## 2 Method

### 2.1 Biophysical model

Forward EMG modeling is traditionally framed as an electric field propagation problem, where a current source density flows through muscle fibers (Farina et al., 2004b,a). Traditional approaches involved solving the quasi-static Maxwell's equations for each fiber at each instant, making the process of simulating realistic forearms with thousands of fibers impractically costly. More recent methods, such as those highlighted in (Maksymenko et al., 2023), adopt a more efficient strategy. These methods solve the equations at a set of static points once, and later integrate fibers as potentials along a subset of these pre-solved points.

Consequently, the forward EMG simulation pipeline is primarily bottlenecked by the initial step of solving equations for a single point source, which incurs the highest computational cost.

#### 2.1.1 Cylindrical model.

For scope of this paper, we use a simplified cylindrical model as our volume conductor seen in 2.1. This model is a standard starting point for simulating EMG and has been extensively studied (Mesin & Farina, 2004; Farina et al., 2004b). It consists of four materials, with isotropic conductivity represented by $\boldsymbol{\sigma}_{iso}$ and muscle fibers characterized by an anisotropic conductivity tensor $\boldsymbol{\sigma}_m$, which aligns with the direction of the fibers.

Our objective is to solve the **parametric volume conduction equation** in this domain for a single point source, with the pinnation angle as a variable parameter. The volume conduction equation is given by (Farina et al., 2004a):

$$\nabla \cdot (\boldsymbol{\sigma}_\phi \nabla v_\phi) = -I \quad \text{in } \boldsymbol{\Omega}, \tag{1}$$
$$\nabla v_\phi \cdot \boldsymbol{n} = 0 \quad \text{on } \partial\boldsymbol{\Omega}\ . \tag{2}$$

Where $\boldsymbol{\Omega}$ is the domain of definition of the three-dimensional volume conductor with boundary $\partial\boldsymbol{\Omega}$, $v_\phi$ is the electric field potential, $\phi$ is the pinnation angle that parameterizes the conductivity tensor $\boldsymbol{\sigma}_\phi$ by altering the direction of the tensor as in Mesin & Farina (2004). Finally $I$ is the input current source density function. The Neumann boundary condition reflect the natural assumption that there is no current flow between the skin and the air.

While this problem can be solved using FEM, the solutions are limited to a single pinnation angle, requiring the construction and solving of a new FEM model for each different angle. Consequently, modeling a continuous range of pinnation angles using FEM becomes impractical.

### 2.2 Data-Driven Solver.

Our proposed solution involves training a neural network (NN) $f(\boldsymbol{x}, \phi; \boldsymbol{\theta}) = v(\boldsymbol{x}; \phi)$ to implicitly represent the continuous range of solutions of the PDE. To ensure the conservation of the prescribed physical relationships, we incorporate PINN or DeepRitz loss (Yu et al., 2018). The surrogate solver can be trained using simulation samples as ground truth or purely by leveraging the physics loss when ground truth data is unavailable. Next, we formally state the problem.

#### 2.2.1 Problem Statement:

Assume a single forcing function $I$, typically representing a single point source, and a range of angles $\phi \in (a, b)$. The range of solutions can be represented by a function with an additional parameter $v(\boldsymbol{x}, \phi) : \boldsymbol{\Omega} \times (a, b) \to \mathbb{R}$.

We then assume that a set of FEM solutions is available for $m$ discrete angles $S_\phi = \{\phi_1, ..., \phi_m\}$. From these solutions, we sample $n$ points for each angle, resulting in a total of $nm$ samples $S_{data} \subset \mathbb{R}^3 \times (a, b)$, which we will refer to as **data points**. We then define the loss function $\mathcal{L}_{data}$ for the training points as follows:

$$\mathcal{L}_{data}(\boldsymbol{\theta}) = \frac{1}{nm} \sum_{j=1}^{m} \sum_{i=1}^{n} ||f(\boldsymbol{x}_i, \phi_j; \boldsymbol{\theta}) - v(\boldsymbol{x}_i, \phi_j)||_2^2. \tag{3}$$

We also sample a set of $n_f$ points in $\mathbb{R}^3$ and $m_f$ angles in $(a, b)$ to construct the set of **collocation points** in $S_{col} \subset \mathbb{R}^3 \times (a, b)$. These collocation points enable the neural network to solve the equation beyond the known values. We define the physics loss using the PINN framework for our PDE as follows:

$$\mathcal{L}_{phys}(\boldsymbol{\theta}) = \frac{1}{n_f m_f} \sum_{j=1}^{m_f} \sum_{i=1}^{n_f} ||\nabla_{\boldsymbol{x}} \cdot (\sigma(\boldsymbol{x}_i, \phi_j) \nabla_{\boldsymbol{x}} f(\boldsymbol{x}_i, \phi_j; \boldsymbol{\theta}) + I(\boldsymbol{x}_i)||_2^2. \tag{4}$$

**Notation:** In this context, we use a slight abuse of notation where $\nabla_{\boldsymbol{x}} \cdot f = f_x + f_y + f_z$ , with $\nabla_{\boldsymbol{x}}$ acting solely on the spatial variables involved in the PDE, and not on the pinnation angle variable $\phi$.

Alternatively the Deep Ritz loss can be used here interchangeably as it achieves a similar purpose. Finally our optimization objective can be defined as follows:

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \lambda \mathcal{L}_{data}(\boldsymbol{\theta}) + \mathcal{L}_{phys}(\boldsymbol{\theta}), \tag{5}$$

Here, $\lambda$ is a weight term that prioritizes fitting to the training data. The next section presents the experimental results of the data-driven solver.

## 3 Experiments

In this section, we provide an overview of our experimental results, focusing on two key questions. First, can a neural network **learn an implicit representation** of the function $v(\boldsymbol{x}, \phi)$ given a training set? Second, how well does this implicit representation **generalize outside the trained region**?

We created a dataset represented by 10 million data points for each angle in $S_\phi = \{0, 5, ...75\}$ using an FEM model. We form three different training sets: $S_\phi = \{0, 5, ...75\}$ , $S_{\phi_8} = \{0, 10, ...70\}$ and $S_{\phi_4} = \{0, 20, ...60\}$ containing solutions for 16, 8, and 4 angles, respectively.

We trained three models using these datasets: **PINN-16**, **PINN-8**, and **PINN-4**. For all models, we used the entire $S_\phi$ as collocation points for the physics-based loss. We validated our models across the entire $S_\phi$. PINN-16 was designed to address the first research question by **fitting the known data**, while PINN-8 and PINN-4 were also **validated on unseen data** to explore generalization.

**Model Architecture Note**: In our experiments we used fully connected networks with residual connection (He et al., 2016) to enhance stability. Our network has width 30 and depth 6.

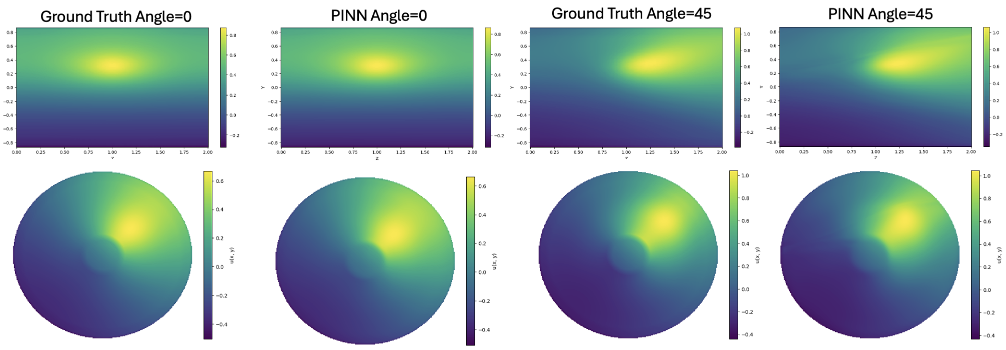


Figure 2: Heatmap of volume conduction solutions at angle=0 (within training region) and angle=45 (outside the region). The top cross section is at $X = 0.5$ and the bottom at $Z = 1.5$.

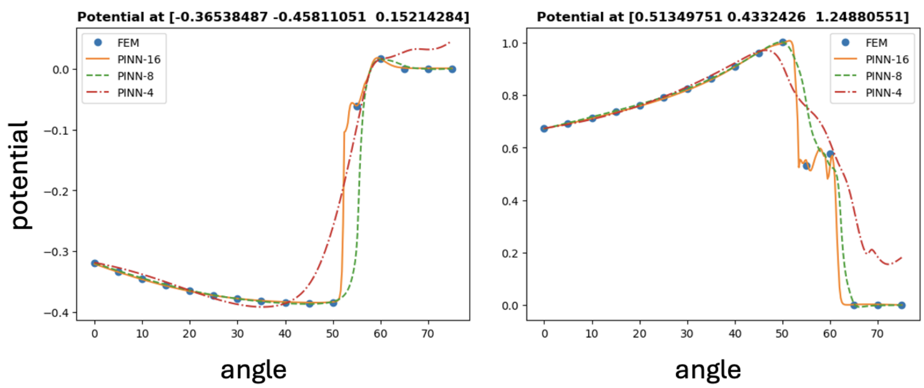


Figure 3: Plot of electric potential of a single point across various angles. The three models are trained in different subsets of the datasets, with PINN-16 on all 16 angles, PINN-8 on half and PINN-4 at only 4 angles.

### 3.1 Preliminary Results

**Overview of the solution:** In Figure 2.2.1, we present two cross-sections of the solutions for different pinnation angles. The flow of potential across the pennate fibers is clearly visible. The fit appears almost identical to the ground truth, excluding the central bone section, where the PINN has difficulty fitting the discontinuous change into the low-conductivity tissue.

**Size and Computational Time Reduction**: The relatively small neural networks, consisting of 5,761 parameters, successfully fit the training data, achieving a mean squared error (MSE) under $10^{-3}$. These results suggest that neural implicit representations are effective for this problem, achieving a 99.995% reduction in model size compared to 16 FEM models, each with 6,944,780 parameters. In terms of computational time, using the evaluation of the solution for 1 million points as a benchmark, the PINN reduced the time by 99.4

**Generalization**: PINN-8 and PINN-4 achieved generalization relative MSE losses of 0.1273 and 0.2577, respectively, across the entire $S_\phi$. Although these errors are significantly higher than the training error, they are not discouraging. With a denser set of angles or more extensive physics-based training, this error can likely be reduced further.

**Conclusion and future work**: The results indicate that data-driven methods can be a highly effective tool for forward EMG modeling, offering significant speedups in simulation and enabling dynamic geometry modeling. Future work will explore additional architectures, such as operator networks Lu et al. (2019), and more complex geometries, such as MRI-based meshes (Maksymenko et al., 2023).

## Acknowledgements

This work was supported by the Imperial-META Wearable Neural Interfaces Research Centre. Additionally, this project received support from the Onassis Foundation under Scholarship ID: F ZT 012-1/2023-2024.

## References

Alexander Kenneth Clarke and Dario Farina. Deep metric learning with locality sensitive mining for self-correcting source separation of neural spiking signals. *IEEE Transactions on Cybernetics*, 2023.

Dario Farina, Luca Mesin, and Simone Martina. Advances in surface electromyographic signal simulation with analytical and numerical descriptions of the volume conductor. *Medical and Biological Engineering and Computing*, 42:467–476, 2004a.

Dario Farina, Luca Mesin, Simone Martina, and Roberto Merletti. A surface emg generation model with multilayer cylindrical description of the volume conductor. *IEEE Transactions on Biomedical Engineering*, 51(3):415–426, 2004b.

Wulfram Gerstner, Henning Sprekeler, and Gustavo Deco. Theory and simulation in neuroscience. *science*, 338(6103):60–65, 2012.

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*, pp. 770–778, 2016.

Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019.

Shihan Ma, Alexander Kenneth Clarke, Kostiantyn Maksymenko, Samuel Deslauriers-Gauthier, Xinjun Sheng, Xiangyang Zhu, and Dario Farina. Conditional generative models for simulation of emg during naturalistic movements. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.

Kostiantyn Maksymenko, Alexander Kenneth Clarke, Irene Mendez Guerra, Samuel Deslauriers-Gauthier, and Dario Farina. A myoelectric digital twin for fast and realistic modelling in deep learning. *Nature Communications*, 14(1):1600, 2023.

Pranav Mamidanna and Dario Farina. Inferring physiological properties of motor neurons using neural posterior estimation. In *ICML 2024 Workshop on Structured Probabilistic Inference* $\{\backslash\&\}$ *Generative Modeling*.

Roberto Merletti, Serge H Roy, Edward Kupa, Silvestro Roatta, and Angelo Granata. Modeling of surface myoelectric signals. ii. model-based signal interpretation. *IEEE Transactions on biomedical engineering*, 46(7):821–829, 1999.

Luca Mesin and Dario Farina. Simulation of surface emg signals generated by muscle tissues with inhomogeneity due to fiber pinnation. *IEEE transactions on biomedical engineering*, 51(9): 1521–1529, 2004.

Diego Pereira Botelho, Kathleen Curran, and Madeleine M Lowery. Anatomically accurate model of emg during index finger flexion and abduction derived from diffusion tensor imaging. *PLoS computational biology*, 15(8):e1007267, 2019.

Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

Justin Sirignano and Konstantinos Spiliopoulos. Dgm: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375:1339–1364, 2018.

Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.

Bing Yu et al. The deep ritz method: a deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.

## A Appendix / supplemental material

Optionally include supplemental material (broader impact, complete proofs, additional experiments and plots) in appendix. All such materials **SHOULD be included in the main submission.**

## D3S3@NeurIPS Paper Checklist (Optional)

The **optional** checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes] , [No] , or [NA] .
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are a part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers.

While "[Yes] " is generally preferable to "[No] ", it is perfectly acceptable to answer "[No] " provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No] " or "[NA] " is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist"**,
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers**.

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: **[TODO]**

   Justification: **[TODO]**

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: **[TODO]**

   Justification: **[TODO]**

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.

- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: **[TODO]**

   Justification: **[TODO]**

   Guidelines:

   - The answer NA means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: **[TODO]**

   Justification: **[TODO]**

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.

- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: **[TODO]**

   Justification: **[TODO]**

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: **[TODO]**

Justification: **[TODO]**

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: **[TODO]**

   Justification: **[TODO]**

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
   - If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: **[TODO]**

   Justification: **[TODO]**

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

Answer: **[TODO]**

Justification: **[TODO]**

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: **[TODO]**

Justification: **[TODO]**

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: **[TODO]**

Justification: **[TODO]**

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.

- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: **[TODO]**

    Justification: **[TODO]**

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: **[TODO]**

    Justification: **[TODO]**

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: **[TODO]**

    Justification: **[TODO]**

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.

- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: **[TODO]**

    Justification: **[TODO]**

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.