Locating Hidden Exoplanets Using Machine Learning

Jason P. Terry
Department of Physics and Astronomy
Center for Simulational Physics
University of Georgia
Athens, GA 30602, USA
jpterry@uga.edu

Sergei Gleyzer
Department of Physics and Astronomy
University of Alabama
Tuscaloosa, AL 35487, USA

Abstract

Exoplanets in protoplanetary disks cause localized deviations from Keplerian velocity in molecular line emission. Current methods of characterizing these deviations are slow and prone to false negatives. We demonstrate that machine learning can quickly and accurately detect planets. We train computer vision models on synthetic observations of protoplanetary disks generated from simulations and apply these models to real observations. The models recreate previous discoveries, accurately locating known planets. A new exoplanet in the disk HD 142666 is identified.

1 Introduction

Kinematic analysis of protoplanetary disks uses molecular line emission to measure the motion of regions within the disk and compares the motion to the default Keplerian motion in which only the star's gravity influences the motion. A planet embedded within the disk will cause localized deviations from Keplerian motion known as "kinks" [Teague et al., 2018, Pinte et al., 2018, 2019]. Identifying these kinks can lead to the identification and characterization of the responsible planet, which improves our understanding of planet formation and disk evolution.

The general process of identifying and verifying planets using kinematic data is slow, difficult, and includes a strong likelihood of false negatives. Of the dozens of observed disks, only a handful have been shown to host a planet via kinematic analysis. paper presents machine learning algorithms applied to protoplanetary disks with the purpose of identifying forming exoplanets more accurately. Our models can locate previously known planets embedded within disks. The models also locate an embedded planet that was previously unreported.

2 Methods

We create 1,000 synthetic protoplanetary disks under various physical and observation conditions (e.g. disk mass, distance, viewing angle, etc.). 13 different parameters are varied in total. The parameter space is sampled using a Latin Hypercube (LHC) [McKay et al., 1979]. We use values from observed ranges inferred from disk surveys, e.g, DSHARP, [Andrews et al., 2009, Andrews et al., 2018, Huang et al., 2018] and widely accepted simulational parameters [Pinte et al., 2018, 2019, 2020]. Disks are assumed to be vertically isothermal.

2.1 Simulations

All 1,000 systems are simulated with 3D smoothed particle hydrodynamics (SPH) using the `PHANTOM` [Price et al., 2018] code (GNU general public license). 1,000,000 particles were used in

Machine Learning and the Physical Sciences Workshop, NeurIPS 2023.

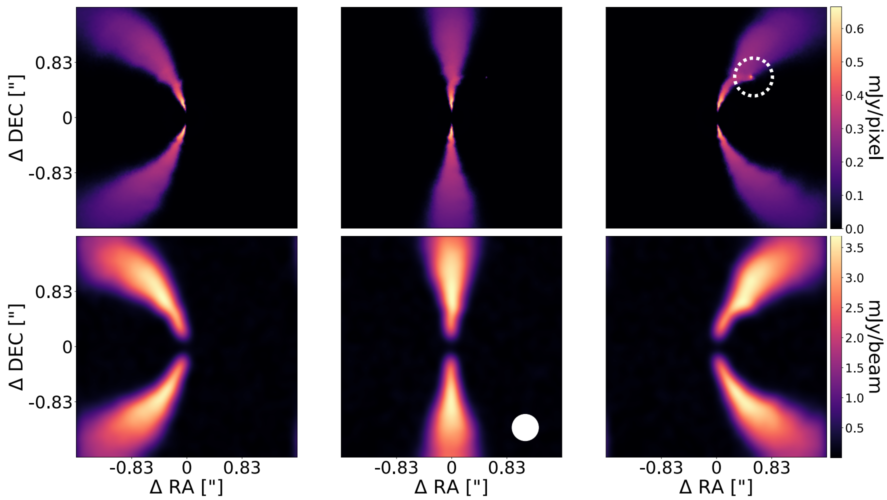

Figure 1: Example raw (top row) and convolved, noisy (bottom row) channel maps in a disk with a planet present. The planet circled by a white dashed line in the upper right panel. The opposite velocity channel is shown in the left column, and the systemic channel is shown in the middle column. The solid white circle in the bottom middle panel is the beam indicating the spatial resolution.

each disk. Data was taken every tenth of an orbit. 25% of the simulations were withheld for testing. 20% of the training data was used for validation.

2.2 Velocity Structure

The simulation results are used to create velocity channel maps in ^{13}CO, a commonly used tracer for H_2 that has been used for the kinematic detection of planets, e.g., [Pinte et al., 2019]. The channel maps are created using the `MCFOST` radiative transfer code [Pinte et al., 2006, 2009] (GNU general public license). The final result for a given velocity channel is an image with 600x600 pixels at a resolution of 1 au/pixel. 10^8 photon packets are used. The velocity channel maps are convolved spatially and spectrally, and noise is added. Figure 1 shows the results of convolving selected velocity channels for a disk with a planet.

2.3 Models

The input for all models is a (600x600xC) image, where C is the number of velocity channels, and the outputs are a classification decision for "no planet" or "at least one planet." We use C = 47, 61, and 75 in order to account for varying resolution in observations. An Adam optimizer is used. The learning rate is updated every 10 epochs by multiplying the current rate by γ. We use cross entropy loss. We train for 50 epochs but allow early stopping with a patience of 8 epochs. 25% of the data was withheld for testing, and 20% of the training data was used for validation.

We use two different models based on PyTorch Torchvision [Paszke et al., 2019] implementations: EfficientNetV2 [Tan and Le, 2021] (ENV2) and RegNet [Xu et al., 2021] (RN). Neither model uses the default hyperparameters or pre-trained weights, which would be under license CC-BY-NC 4.0. We perform Bayesian hyperparameter tuning using `WANDB` [Biewald, 2020] to find separate sets of hyperparameters that minimize the validation loss. The initial hyperparameters for these sweeps are based on default versions “EfficientNetV2 S” and “RegNetY 16GF.” Table 1 gives hyperparameters as determined by the Bayesian hyperparameter sweep. The models were too large to fit on available GPUs, so they were all trained on CPUs.

Model	Parameters (Millions)	Learning Rate	γ	Depth	Dropout	Training Time (hours)
ENV2 (47)	20.2	0.0102	0.0429	-	0.0391	3
ENV2 (61)	20.2	0.0087	0.0326	-	0.0278	7
ENV2 (75)	20.2	0.0084	0.0298	-	0.0335	13
RN (47)	51.0	0.0368	0.0087	14	-	19
RN (61)	62.8	0.0012	0.0428	15	-	16
RN (75)	114	0.0033	0.0176	21	-	22

Table 1: A subset of hyperparameters for all models as determined by the Bayesian sweep using `WANDB`.

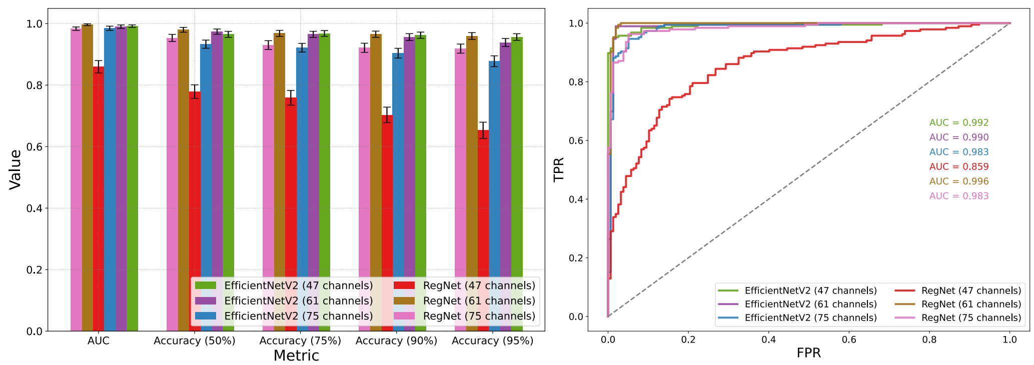

Figure 2: Left: Various metrics calculated from the withheld test set for all models. Right: Corresponding ROC curves. Error bars show 95% confidence intervals that are calculated by bootstrapping each metric 1,000 times using random selections of 80% of the test data. The percentages next to the accuracy labels denote the decision threshold.

3 Results and Discussion

3.1 Machine learning Models

Figure 2 gives several metrics for model performance and the ROC curves. We consider several values of the softmax output for the classification decision threshold. It is encouraging that, for all other models than the RN (47), there are no qualitative changes in the results based on the decision threshold.

3.2 Observations

We demonstrate the effectiveness of our model by applying it to real telescope observations. HD 97048 hosts a planet that was discovered by Pinte et al. [2019] using kinematic observations. Using the same data as Pinte et al. [2019] (ADS/JAO.ALMA#2016.1.00826.S), we show that our models replicate the prediction and estimated location of the forming exoplanet (Figure 3). Passing this essential test gives support to the idea that the models can at least match human performance. Similar results were found for HD 163296, which is also known to host a planet [Teague et al., 2018, Pinte et al., 2018].

3.2.1 HD 142666

Recreating previous discoveries is encouraging, but the purpose of using machine learning is to facilitate new discoveries. To do so, we apply our models to all disks in the DSHARP catalogue [Andrews et al., 2018] that do not have reported planets. The results for HD 142666 (ADS/JAO.ALMA #2016.1.00484.L) show unambiguous signatures of an embedded planet. This disk has been analyzed for years, yet no human found the planet. In less than five minutes, all six of our models strongly identified and localized the planet. Figure 4 shows these results. The upper row shows the velocity channels in which the planet is present. The lower row shows the internal activations of three different

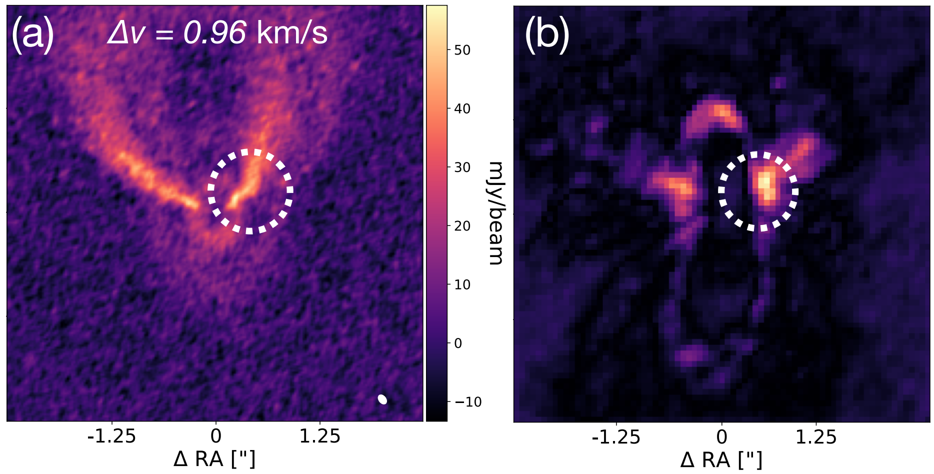

Figure 3: (a): $\Delta v = 0.96$ km/s channel. (b): Example activation structure using RN (47). The estimated planet location ([Pinte et al., 2019]) is given inside the white circle.

models, which clearly highlight those same velocity channels. Follow-up simulations confirmed the planet by existing standards Teague et al. [2018], Pinte et al. [2018, 2019].

3.3 Limitations

Our work comes with an important limitation regarding the generated datasets: all simulations are run with SPH. Grid-based codes could be used to create all of the data or supplement the SPH data in order to span algorithmic differences. For the sake of consistency and simplicity of interpretation, we did not do this. Supplementation with grid-based data could be a good direction for future work. It is possible that the method of domain adaptation [Ben-David et al., 2010], which has already found success in astronomy [Vilalta et al., 2019, Alexander et al., 2023, Ćiprijanović et al., 2022], could be used to encourage the models to overcome any differences between grid-based and SPH datasets.

4 Conclusions

Our results show that a machine learning model trained on synthetic data can determine the presence and location of forming exoplanets in real telescope observations of protoplanetary accretion disks. We apply this model to the observed HD 97048 and HD 142666 disks. The presence and location of the planet in HD 97048 is corroborated with confidences of up to $> 99\%$ and activations highlight the same location given by Pinte et al. [2019]. The models identify a previously unreported exoplanet embedded within HD 142666. These results give a strong endorsement for the use of machine learning in observational astronomy. Models can match and exceed human proficiency in a fraction of the time, thereby facilitating discovery and analysis. This will be of utmost importance as both ALMA and JWST continue to deliver larger and larger disk survey datasets and next generation telescopes —such as ngVLA and the SKA —come online.

Acknowledgments and Disclosure of Funding

This paper makes use of the following ALMA data: ADS/JAO.ALMA#2016.1.00826.S , ADS/JAO.ALMA#2013.1.00601.S, and ADS/JAO.ALMA #2016.1.00484.L. ALMA is a partnership of ESO (representing its member states), NSF (USA) and NINS (Japan), together with NRC (Canada), MOST and ASIAA (Taiwan), and KASI (Republic of Korea), in cooperation with the Republic of Chile. The Joint ALMA Observatory is operated by ESO, AUI/NRAO and NAOJ. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. SPH results are visualized using `SPLASH` [Price, 2007]. JT was a participant in the 2022 and 2023 ML4SCI Google Summer of Code programs. This study

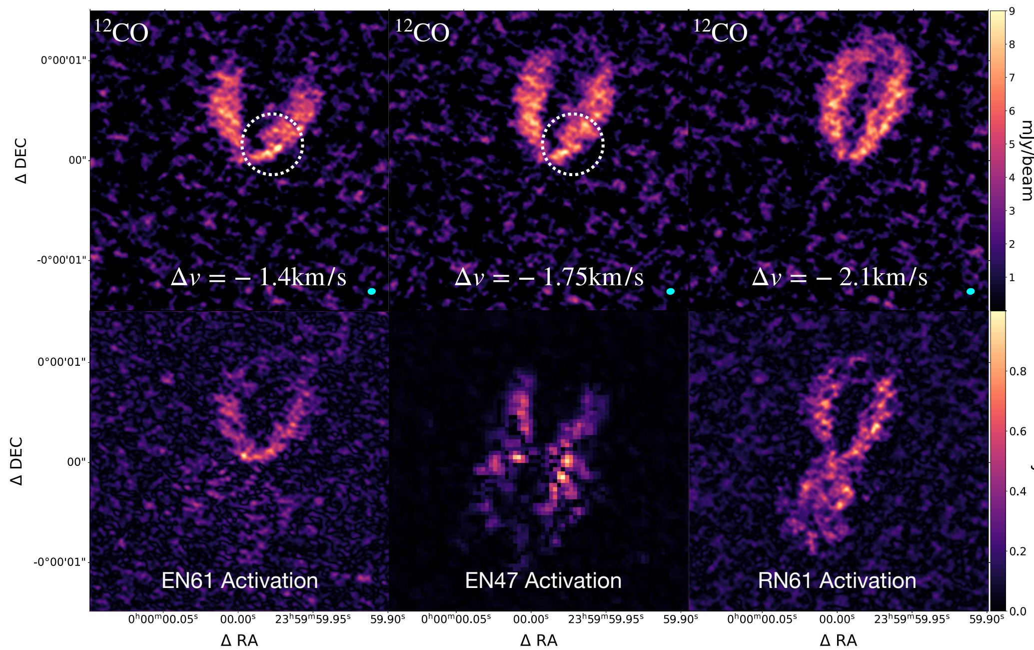

Figure 4: HD 142666 structure (^{12}CO: $J = 2 \to 1$) and activations. Upper left: $\Delta v = -1.4$ km/s channel with kink circled in white. Upper middle: $\Delta v = -1.75$ channel with kink circled in white. Upper right: $\Delta v = -2.1$ channel. Bottom row: selected activations that roughly correspond to the channels in the upper row.

was supported in part by resources and technical expertise from the Georgia Advanced Computing Resource Center, a partnership between the University of Georgia's Office of the Vice President for Research and Office of the Vice President for Information Technology. SG was supported in part by U.S. National Science Foundation Award No. 2108645.

5 Impact Statement

We do not anticipate any negative societal effects from our work. There is little damage that can be done by claiming a planet is present. This is a strong claim that requires multiple steps of verification, so the main issue a false positive would create is using the time of experts. Should the model consistently miss planets, that would be unfortunate. However, there are teams pouring over this data with the same objective and gaining no useful insight from our work is simply the status quo. On the other hand, if our work is consistently successful, it could be important for speeding up the analysis of protoplanetary disks, which would facilitate the progress of planet formation theory.

References

Stephon Alexander, Sergei Gleyzer, Hanna Parul, Pranath Reddy, Marcos Tidball, and Michael W. Toomey. Domain Adaptation for Simulation-based Dark Matter Searches with Strong Gravitational Lensing. *ApJ*, 954(1):28, September 2023. doi: 10.3847/1538-4357/acdfc7.

Sean M. Andrews, D. J. Wilner, A. M. Hughes, Chunhua Qi, and C. P. Dullemond. PROTOPLANETARY DISK STRUCTURES IN OPHIUCHUS. *The Astrophysical Journal*, 700(2):1502–1523, jul 2009. doi: 10.1088/0004-637x/700/2/1502. URL `https://doi.org/10.1088/0004-637x/700/2/1502`.

Sean M. Andrews, Jane Huang, Laura M. Pérez, Andrea Isella, Cornelis P. Dullemond, Nicolás T. Kurtovic, Viviana V. Guzmán, John M. Carpenter, David J. Wilner, Shangjia Zhang, Zhaohuan

Zhu, Tilman Birnstiel, Xue-Ning Bai, Myriam Benisty, A. Meredith Hughes, Karin I. Öberg, and Luca Ricci. The Disk Substructures at High Angular Resolution Project (DSHARP). I. Motivation, Sample, Calibration, and Overview. *ApJL*, 869(2):L41, December 2018. doi: 10.3847/2041-8213/aaf741.

Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79(1):151–175, 2010. doi: 10.1007/s10994-009-5152-4. URL `https://doi.org/10.1007/s10994-009-5152-4`.

Lukas Biewald. Experiment tracking with weights and biases, 2020. URL `https://www.wandb.com/`. Software available from wandb.com.

Aleksandra Ćiprijanović, Diana Kafkes, Gregory Snyder, F. Javier Sánchez, Gabriel Nathan Perdue, Kevin Pedro, Brian Nord, Sandeep Madireddy, and Stefan M. Wild. DeepAdversaries: examining the robustness of deep learning models for galaxy morphology classification. *Machine Learning: Science and Technology*, 3(3):035007, September 2022. doi: 10.1088/2632-2153/ac7f1a.

Jane Huang, Sean M. Andrews, Cornelis P. Dullemond, Andrea Isella, Laura M. Pérez, Viviana V. Guzmán, Karin I. Öberg, Zhaohuan Zhu, Shangjia Zhang, Xue-Ning Bai, Myriam Benisty, Tilman Birnstiel, John M. Carpenter, A. Meredith Hughes, Luca Ricci, Erik Weaver, and David J. Wilner. The Disk Substructures at High Angular Resolution Project (DSHARP). II. Characteristics of Annular Substructures. *ApJL*, 869(2):L42, December 2018. doi: 10.3847/2041-8213/aaf740.

M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2): 239–245, 1979. ISSN 00401706. URL `http://www.jstor.org/stable/1268522`.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *arXiv e-prints*, art. arXiv:1912.01703, December 2019.

C. Pinte, F. Ménard, G. Duchêne, and P. Bastien. Monte Carlo radiative transfer in protoplanetary disks. , 459(3):797–804, December 2006. doi: 10.1051/0004-6361:20053275.

C. Pinte, T. J. Harries, M. Min, A. M. Watson, C. P. Dullemond, P. Woitke, F. Ménard, and M. C. Durán-Rojas. Benchmark problems for continuum radiative transfer. High optical depths, anisotropic scattering, and polarisation. , 498(3):967–980, May 2009. doi: 10.1051/0004-6361/200811555.

C. Pinte, D. J. Price, F. Ménard, G. Duchêne, W. R. F. Dent, T. Hill, I. de Gregorio-Monsalvo, A. Hales, and D. Mentiplay. Kinematic Evidence for an Embedded Protoplanet in a Circumstellar Disk. *ApJL*, 860(1):L13, June 2018. doi: 10.3847/2041-8213/aac6dc.

C. Pinte, G. van der Plas, F. Ménard, D. J. Price, V. Christiaens, T. Hill, D. Mentiplay, C. Ginski, E. Choquet, Y. Boehler, G. Duchêne, S. Perez, and S. Casassus. Kinematic detection of a planet carving a gap in a protoplanetary disk. *Nature Astronomy*, 3:1109–1114, August 2019. doi: 10.1038/s41550-019-0852-6.

C. Pinte, D. J. Price, F. Ménard, G. Duchêne, V. Christiaens, S. M. Andrews, J. Huang, T. Hill, G. van der Plas, L. M. Perez, A. Isella, Y. Boehler, W. R. F. Dent, D. Mentiplay, and R. A. Loomis. Nine Localized Deviations from Keplerian Rotation in the DSHARP Circumstellar Disks: Kinematic Evidence for Protoplanets Carving the Gaps. *ApJL*, 890(1):L9, February 2020. doi: 10.3847/2041-8213/ab6dda.

Daniel J. Price. splash: An Interactive Visualisation Tool for Smoothed Particle Hydrodynamics Simulations. , 24(3):159–173, October 2007. doi: 10.1071/AS07022.

Daniel J. Price, James Wurster, Terrence S. Tricco, Chris Nixon, Stéven Toupin, Alex Pettitt, Conrad Chan, Daniel Mentiplay, Guillaume Laibe, Simon Glover, Clare Dobbs, Rebecca Nealon, David Liptai, Hauke Worpel, Clément Bonnerot, Giovanni Dipierro, Giulia Ballabio, Enrico Ragusa, Christoph Federrath, Roberto Iaconi, Thomas Reichardt, Duncan Forgan, Mark Hutchison, Thomas

Constantino, Ben Ayliffe, Kieran Hirsh, and Giuseppe Lodato. Phantom: A Smoothed Particle Hydrodynamics and Magnetohydrodynamics Code for Astrophysics. , 35:e031, September 2018. doi: 10.1017/pasa.2018.25.

Mingxing Tan and Quoc V. Le. EfficientNetV2: Smaller Models and Faster Training. *arXiv e-prints*, art. arXiv:2104.00298, April 2021.

Richard Teague, Jaehan Bae, Edwin A. Bergin, Tilman Birnstiel, and Daniel Foreman-Mackey. A Kinematical Detection of Two Embedded Jupiter-mass Planets in HD 163296. *ApJL*, 860(1):L12, June 2018. doi: 10.3847/2041-8213/aac6d7.

Ricardo Vilalta, Kinjal Dhar Gupta, Dainis Boumber, and Mikhail M. Meskhi. A General Approach to Domain Adaptation with Applications in Astronomy. , 131(1004):108008, October 2019. doi: 10.1088/1538-3873/aaf1fc.

Jing Xu, Yu Pan, Xinglin Pan, Steven Hoi, Zhang Yi, and Zenglin Xu. RegNet: Self-Regulated Network for Image Classification. *arXiv e-prints*, art. arXiv:2101.00590, January 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? We demonstrate that machine learning can locate exoplanets within protoplanetary disks using both simulated and observational data.
   (b) Did you describe the limitations of your work? Section 3.3 is dedicated to discussing the limitations of our data and methods.
   (c) Did you discuss any potential negative societal impacts of your work? Section 5 addresses this topic.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? Physical assumptions and parameters are given in Section 2
   (b) Did you include complete proofs of all theoretical results? We used no proofs.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? All special packages used for (i.e. excluding things such as NumPy) are cited and linked.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? Sections 2.2 and 2.3 address this.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? Errors are reported in Figure 2
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? Section 2.3 addresses this.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? All codes are cited and linked.
   (b) Did you mention the license of the assets? Licenses are described after the packages are introduced.
   (c) Did you include any new assets either in the supplemental material or as a URL? The only new asset is a private GitHub repository that will not be made public until this project has concluded. It will be under the Creative Commons license.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? We use our own data.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Our data includes nothing remotely related to such topics.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? No use of crowdsourcing or humans.
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? No use of crowdsourcing or humans.
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? No use of crowdsourcing or humans.