
Self-supervised similarity search for large scientific datasets

George Stein

Berkeley Center for Cosmological Physics
Lawrence Berkeley National Laboratory
Berkeley, CA 94720
gstein@berkeley.edu

Peter Harrington

Lawrence Berkeley National Laboratory
Berkeley, CA 94720
pharrington@lbl.gov

Jacqueline Blaum

Department of Astronomy, University of California, Berkeley
Lawrence Berkeley National Laboratory
Berkeley, CA 94720
jrblaum@berkeley.edu

Tomislav Medan

Lawrence Berkeley National Laboratory
Berkeley, CA 94720
tmedan@lbl.gov

Zarija Lukić

Lawrence Berkeley National Laboratory
Berkeley, CA 94720
zarija@lbl.gov

Abstract

We present the use of self-supervised learning to explore and exploit large unlabeled datasets. Focusing on 42 million galaxy images from the latest data release of the Dark Energy Spectroscopic Instrument (DESI) Legacy Imaging Surveys, we first train a self-supervised model to distill low-dimensional representations that are robust to symmetries, uncertainties, and noise in each image. We then use the representations to construct and publicly release an interactive semantic similarity search tool. We demonstrate how our tool can be used to rapidly discover rare objects given only a single example, increase the speed of crowd-sourcing campaigns, and construct and improve training sets for supervised applications. While we focus on images from sky surveys, the technique is straightforward to apply to any scientific dataset of any dimensionality. The similarity search web app can be found at github.com/georgestein/galaxy_search.

1 Introduction

Scientific datasets continue to increase in both size and complexity, yet often lack high-quality labels that can be used to explore and understand the natural distribution of objects and features within them. As such, we require accessible automated tools to select sets of interesting objects and to assess the rarity of individual observations. One such important tool for exploration and data discovery is similarity search, which, given an input sample, rapidly returns semantically similar instances. Also known as query by example, content-based instance retrieval, or reverse instance search, this can be used for a number of scientific investigations including discovering rare objects given only a single example, flagging and identifying bad data, and rapidly constructing and improving training sets for supervised applications.

Quantifying the similarity between two data samples is difficult in the measurement space, as samples that contain the same scientific information can have significantly different data values. For example, rotating an image will likely drastically change individual pixel values, yet does not change the semantic information in that image (assuming rotational symmetry). Therefore, instance-based retrieval methods first extract a lower-dimensional representation from each data sample, and then construct a notion of distance or similarity in this feature space. While commonly used feature extraction methods such as t-SNE [17], UMAP [12], pre-trained supervised networks [15], or autoencoders and their variants [10, 11], have been shown to extract relevant information for a number of tasks, they can also put focus on scientifically unimportant features or noise. In the case of pre-trained supervised networks, which rely on human-annotated labels, data samples can end up being grouped more by their predetermined (human-constructed) classes than purely by their semantic similarity. The key for extracting scientifically useful features is to explicitly inform the model of the symmetries, uncertainties, or instrumental noise in the dataset, so that the resulting features are robust to these perturbations.

In this vein, contrastive self-supervised learning has achieved remarkable results in recent years, extracting features (from potentially vast uncurated datasets, as in [7]) which have remarkable predictive power and semantic structure despite never using any supervision labels during training [2, 4, 8, 3]. Contrastive learning allows one to easily tailor the training process such that extracted features exclude undesired or uninformative properties, like symmetries or observational noise in a given dataset, by adjusting the data augmentations used. Thus, the features extracted from contrastive self-supervised models are an excellent choice for use in scientific similarity searches. Here we present the use of self-supervised learning to distill low-dimensional representations from 42 million galaxy images from the latest data release of the Dark Energy Spectroscopic Instrument (DESI) Legacy Imaging Surveys, and then construct and publicly release a semantic similarity search tool that can instantly retrieve the most similar images to any given query. We show a schematic of our approach in Figure 1. While we focus on an example from sky surveys, the technique we present can in principle be used on any scientific dataset of any dimensionality.

2 Methods

Data Our dataset consists of galaxy images from the latest data release of the DESI Legacy Imaging Surveys¹, Data Release 9 (DR9), released in January of 2021 [6]. Our self-supervised approach allows us to learn from the entire dataset and does not require any labels or restrictive data cuts, so we construct a dataset of the brightest galaxies from the south region of DR9, as this region has the lowest level of noise. We focused on the 42,272,646 galaxies with a z-band magnitude < 20 , as galaxies dimmer than this threshold generally have an angular size that only covers a few pixels of the image. Centered on each galaxy we extracted a 152×152 pixel cutout in the three optical bands (g, r, z) at the 0.262 arcsecond resolution of the telescope. Our network was trained on 96×96 pixel crops, which is sufficient to cover the angular extent of all but the most near-by galaxies, but the 152^2 pixel versions are required in order to perform the required set of data augmentations during training. This results in a dataset size of 10TB.

Self-supervised learning We closely follow [9] in designing the architecture and training procedure for our self-supervised model, which is based on MoCov2 [4]. In this setting, the backbone of the model is a CNN encoder that takes an image \mathbf{x} as input and produces a lower dimensional representation \mathbf{z} . The encoder learns to make meaningful representations by associating augmented views of the same image as similar, and views of different images as dissimilar, via a contrastive loss function. We use the same ResNet50 network and training hyperparameters as [9], but increase the queue length to $K = 262,144$ to accommodate our larger training set. We choose the following set of augmentations, applying each of them in succession to images during pre-training with the order listed below [see 16, for motivation]:

- *Galactic extinction:* We de-redden the image according to its tabulated SFD $E(B - V)$ value [14], then randomly sample a new $E(B - V)$ value from a log-normal distribution fit to the dataset.
- *Rotation/Orientation:* We randomly flip the image across each axis with 50% probability, then rotate by a random angle sampled from $\mathcal{U}(0, 2\pi)$.

¹legacysurvey.org/

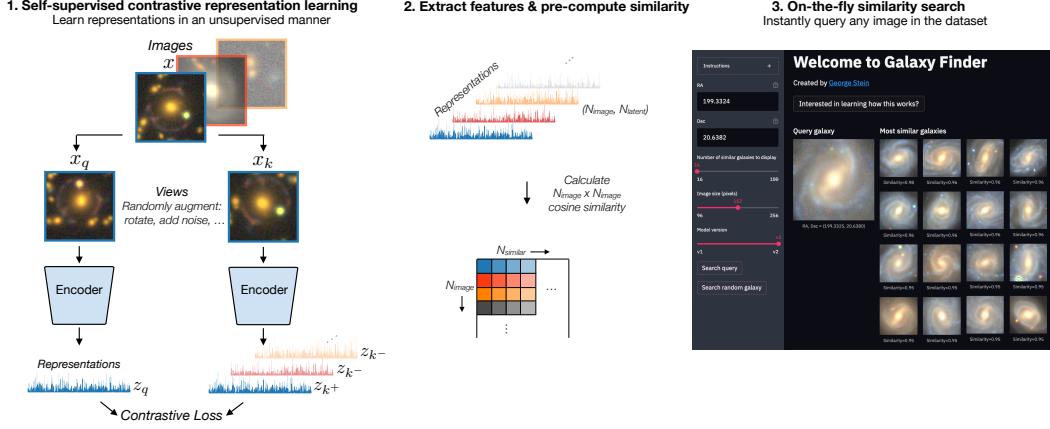


Figure 1: Contrastive self-supervised learning to extract representations which are robust to symmetries, uncertainties, and instrumental noise in the dataset (left), feature extraction and similarity computation (center), and our public interactive similarity search tool (right) – github.com/georgestein/galaxy_search.

- *Size Scaling:* We randomly re-size the image to between 90% and 110% of its original size by rescaling with bilinear interpolation.
- *Point Spread Function (PSF) blur:* We sample a Gaussian blur for each channel by sampling from log-normal fits to the PSF distribution of the data.
- *Jitter and crop:* We translate the image center along each respective axis by $\mathcal{U}(-7, 7)$ pixels, before cropping out the central 96×96 pixels.
- *Gaussian noise:* We sample a Gaussian noise level from log-normal distributions fit to the noise distribution in each filter channel.

With these augmentations, we perform self-supervised pretraining on a curated subset of 3.5 million galaxies sampled uniformly by z -band magnitude. Including all objects was deemed to be not worth the order of magnitude more computation time, as the majority of the fainter galaxies have a spatial size in the image spanning no more than a few pixels. We are able to train the model in just 8 hours on 8 NVIDIA Tesla V100 GPUs. Then, we apply the CNN encoder to all images in our dataset, extracting a representation vector \mathbf{z} of length 2048 for each galaxy.

Similarity Search Self-supervised pre-training provides a robust measure of semantic similarity between any set of images. The representations extracted from each image in the dataset retain the overall information such as the size and relative orientation of the galaxies, the number of sources, the clustering, and the color, while removing the noise and symmetries described by the data augmentations [for a visualization of how galaxy properties are organized in the representation space of a self-supervised model, see 9].

The contrastive loss in self-supervised pre-training encourages semantically similar images to have similar representations, so we can easily define a similarity metric in representation space (which is a lower dimensionality than the input images). We choose to measure similarity between any two images \mathbf{x}_i and \mathbf{x}_j by the cosine similarity (normalized scalar product) of their representation vectors \mathbf{z}_i and \mathbf{z}_j :

$$\text{similarity}(\mathbf{z}_i, \mathbf{z}_j) = \mathbf{z}_i \cdot \mathbf{z}_j / (\|\mathbf{z}_i\| \|\mathbf{z}_j\|). \quad (1)$$

To speed up similarity search we pre-calculate the 999 most similar images to each of the 42 million images, and save the resulting “similarity array” (size $\sim 42\text{M} \times 1,000$) for rapid querying. Computations are performed on 8 GPUs using Facebook AI Research’s Faiss library for efficient similarity search.² To avoid unnecessary computations we did not compute the full $N \times N$ similarity matrix ($N = 42$ million), but for each source limited the search to sources within a magnitude range

²github.com/facebookresearch/faiss. While we calculate similarity without employing any of the compression techniques available in Faiss we found that it still outperformed its NumPy equivalent on CPU, and benefited from an additional speed up on GPU




Figure 2: Query examples (left-most panels) and a sample of five similar galaxies returned through self-supervised similarity search, showing ring galaxies (top row), spiral galaxies (second row), strong gravitational lenses (third row) [16], and galaxy mergers and images corrupted by nearby bright stars (bottom row).

of 0.5 in r-band, as we found that this resulted in nearly equivalent results as the full search due to the model naturally learning that brighter galaxies do not have significant similarity with dimmer ones.

3 Results

We first experimented with which representation layer to use for optimal similarity computation. The layer before the projection head, rather than the final layer where the contrastive loss is computed, has been shown to improve linear classification results on ImageNet by $>10\%$ [3], but this does not necessarily directly transfer to an improved similarity search. We construct separate similarity search models using both sets of representations (2048 dimensional and 128 dimensional representations, respectively), and augment a large random sample of images with the same transformations used during training. Performing similarity search with the original images as a query, we find that the augmented versions are nearly always returned as the most similar sample to the original image for both models, with little to no quantitative difference. While perhaps unsurprising given the clear semantic similarity of query examples and the most similar images returned, this test provides evidence that both representation layers of our self-supervised model are robust to the input perturbations, as designed. Whether this is a result of our augmentations being an incomplete description of the full uncertainty in the dataset is left for future work. From visual inspections of a number of similarity searches we conclude that the most similar images returned by the pre-projection (2048-dimensional) representations marginally outperform those from the final post-projection representations.

In Figure 2 we provide a few example searches using astronomical objects of interest, and show similar galaxies hand-selected from the top ~ 10 returned. As gravitational lenses are exceptionally rare, for these two panels we instead examine the top few hundred examples returned and select the most interesting. We find remarkable results for a large diversity of objects – the semantic similarity as measured by our self-supervised model goes far beyond what is possible by hand-selected color or magnitude cuts or hand-constructed filters. It is clear that the galaxies returned for each query do not simply have similar orientations or noise levels – rather, the notion of similarity in the self-supervised representations extends far beyond such basic characteristics, and reflects the rich set of object categories in the data.

For lack of space we highlight only two examples: 1.) Strong gravitational lens finding poses a significant challenge, as they are amongst the rarest objects in the universe (~ 1 in $\mathcal{O}(10^4)$ [5]). Finding them requires large visual inspection efforts with standard supervised learning techniques due to limitations in the small number of labels available for training [see 13, and references within], yet here we queried for a single serendipitously discovered lens [1] (third row, on left) and instantly discovered many additional lenses. In [16] we utilize similarity search alongside supervised classification to discover $\sim 1,200$ new strong lenses. 2.) Ongoing crowd sourcing campaigns to

discover and classify objects such as ring-galaxies³ can be massively sped up through self-supervised similarity search by pre-selecting unlabelled samples based on minimal similarity searches, and then inspecting them in parallel.

Similarity search on self-supervised representations may focus on any features that are not explicitly excluded by the data augmentations. For example, an input query to our search tool that contains bright foreground stars or galaxies can have the “similarity focus” drawn away from object of interest, and will instead return other images with similar foreground features. In this way this technique is limited by the augmentations chosen during training. A potential workaround is to include such foreground effects as data augmentations in self-supervised pre-training — for example by randomly adding bright foreground objects to a given view of the original image — which would incentivize the representations to be less susceptible to these perturbations. We leave this for future investigation. Additionally, for this work we chose a minimal set of task-agnostic augmentations to use during training, such that the representations are directly useful for a broad number of survey science tasks. If the tool was instead desired for more specific tasks, for example an emphasis on color over shape, the strength of the augmentations, or the inclusion of additional ones, can help shape the similarity criteria.

4 Conclusion

We have demonstrated that a self-supervised learning model can be used to extract information-rich representations from unlabelled data that are robust to symmetries, uncertainties, or instrumental noise in the dataset. These representations provide a direct quantification of similarity between any two data samples, an extremely useful ability given that most real-world scientific datasets have no objective measure of “true” similarity between samples. As the model was trained only to distinguish the natural distribution of features within the data sample, it is not limited to identification of a specific class of object, but can be used to facilitate discovery of any object(s) that exist in the dataset.

5 Broader impact

Generalized self-supervised models greatly reduce the barrier to entry for working with modern datasets, and open up a number of collaborative avenues that were previously unavailable. Rather than each team working alone to perform classification or regression tasks on large scientific datasets – applying for time on large GPU computing systems, downloading massive datasets, learning to train models in parallel, etc. – the initial model training can be undertaken separately, and the fully trained model and image/representation pairs can be shared to allow for rapid scientific investigations that are not hindered by computational complexity or resource allocation. This creates the potential to further democratise scientific investigations on large datasets to institutions that do not have the same computational accessibility as tier-one research institutes or large technology companies.

6 Acknowledgements

We would like to thank Rongpu Zhou for his significant help on using the DESI Legacy Survey data, Dustin Lang for providing access to the image-cutout service at NERSC, and Md Abul Hayat and Mustafa Mustafa for their pioneering efforts on self-supervised learning for sky surveys.

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. This research also used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. G.S., J.B., T.M., and Z.L. were partially supported by the DOE’s Office of Advanced Scientific Computing Research and Office of High Energy Physics through the Scientific Discovery through Advanced Computing (SciDAC) program.

³see the ongoing crowd-sourced ring galaxy campaign blog.galaxyzoo.org/2021/09/23/new-galaxy-zoo-mobile-challenge-ring-galaxies/

References

- [1] E. J. Buckley-Geer, H. Lin, E. R. Drabek, S. S. Allam, D. L. Tucker, R. Armstrong, W. A. Barkhouse, E. Bertin, M. Brodwin, S. Desai, J. A. Frieman, S. M. Hansen, F. W. High, J. J. Mohr, Y. T. Lin, C. C. Ngeow, A. Rest, R. C. Smith, J. Song, and A. Zenteno. The Serendipitous Observation of a Gravitationally Lensed Galaxy at $z = 0.9057$ from the Blanco Cosmology Survey: The Elliot Arc. *The Astrophysical Journal*, 742(1):48, Nov. 2011. doi: 10.1088/0004-637X/742/1/48.
- [2] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin. Emerging Properties in Self-Supervised Vision Transformers. *arXiv e-prints*, art. arXiv:2104.14294, Apr. 2021.
- [3] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A Simple Framework for Contrastive Learning of Visual Representations. *arXiv e-prints*, art. arXiv:2002.05709, Feb. 2020.
- [4] X. Chen, H. Fan, R. Girshick, and K. He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [5] T. E. Collett. The Population of Galaxy-Galaxy Strong Lenses in Forthcoming Optical Imaging Surveys. *The Astrophysical Journal*, 811(1):20, Sept. 2015. doi: 10.1088/0004-637X/811/1/20.
- [6] A. Dey, D. J. Schlegel, D. Lang, R. Blum, K. Burleigh, X. Fan, J. R. Findlay, D. Finkbeiner, D. Herrera, S. Juneau, M. Landriau, M. Levi, I. McGreer, A. Meisner, A. D. Myers, J. Moustakas, P. Nugent, A. Patej, E. F. Schlafly, A. R. Walker, F. Valdes, B. A. Weaver, C. Yèche, H. Zou, X. Zhou, B. Abareshi, T. M. C. Abbott, B. Abolfathi, C. Aguilera, S. Alam, L. Allen, A. Alvarez, J. Annis, B. Ansarinejad, M. Aubert, J. Beechert, E. F. Bell, S. Y. BenZvi, F. Beutler, R. M. Bielby, A. S. Bolton, C. Briceño, E. J. Buckley-Geer, K. Butler, A. Calamida, R. G. Carlberg, P. Carter, R. Casas, F. J. Castander, Y. Choi, J. Comparat, E. Cukanovaite, T. Delubac, K. DeVries, S. Dey, G. Dhungana, M. Dickinson, Z. Ding, J. B. Donaldson, Y. Duan, C. J. Duckworth, S. Eftekharzadeh, D. J. Eisenstein, T. Etourneau, P. A. Fagrelus, J. Farihi, M. Fitzpatrick, A. Font-Ribera, L. Fulmer, B. T. Gänsicke, E. Gaztanaga, K. George, D. W. Gerdes, S. G. A. Gontcho, C. Gorgoni, G. Green, J. Guy, D. Harmer, M. Hernandez, K. Honscheid, L. W. Huang, D. J. James, B. T. Jannuzi, L. Jiang, R. Joyce, A. Karcher, S. Karkar, R. Kehoe, J.-P. Kneib, A. Kueter-Young, T.-W. Lan, T. R. Lauer, L. Le Guillou, A. Le Van Suu, J. H. Lee, M. Lesser, L. Perreault Levasseur, T. S. Li, J. L. Mann, R. Marshall, C. E. Martínez-Vázquez, P. Martini, H. du Mas des Bourboux, S. McManus, T. G. Meier, B. Ménard, N. Metcalfe, A. Muñoz-Gutiérrez, J. Najita, K. Napier, G. Narayan, J. A. Newman, J. Nie, B. Nord, D. J. Norman, K. A. G. Olsen, A. Paat, N. Palanque-Delabrouille, X. Peng, C. L. Poppett, M. R. Poremba, A. Prakash, D. Rabinowitz, A. Raichoor, M. Rezaie, A. N. Robertson, N. A. Roe, A. J. Ross, N. P. Ross, G. Rudnick, S. Safonova, A. Saha, F. J. Sánchez, E. Savary, H. Schweiker, A. Scott, H.-J. Seo, H. Shan, D. R. Silva, Z. Slepian, C. Soto, D. Sprayberry, R. Staten, C. M. Stillman, R. J. Stupak, D. L. Summers, S. Sien Tie, H. Tirado, M. Vargas-Magaña, A. K. Vivas, R. H. Wechsler, D. Williams, J. Yang, Q. Yang, T. Yapici, D. Zaritsky, A. Zenteno, K. Zhang, T. Zhang, R. Zhou, and Z. Zhou. Overview of the DESI Legacy Imaging Surveys. *The Astronomical Journal*, 157(5):168, May 2019. doi: 10.3847/1538-3881/ab089d.
- [7] P. Goyal, M. Caron, B. Lefaudeaux, M. Xu, P. Wang, V. Pai, M. Singh, V. Liptchinsky, I. Misra, A. Joulin, and P. Bojanowski. Self-supervised Pretraining of Visual Features in the Wild. *arXiv e-prints*, art. arXiv:2103.01988, Mar. 2021.
- [8] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. D. Guo, M. Gheshlaghi Azar, B. Piot, K. Kavukcuoglu, R. Munos, and M. Valko. Bootstrap your own latent: A new approach to self-supervised Learning. *arXiv e-prints*, art. arXiv:2006.07733, June 2020.
- [9] M. A. Hayat, G. Stein, P. Harrington, Z. Lukić, and M. Mustafa. Self-supervised representation learning for astronomical images. *The Astrophysical Journal Letters*, 911(2):L33, apr 2021. doi: 10.3847/2041-8213/abf2c7. URL <https://doi.org/10.3847/2041-8213/abf2c7>.
- [10] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006. doi: 10.1126/science.1127647. URL <https://www.science.org/doi/abs/10.1126/science.1127647>.

- [11] D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. *arXiv e-prints*, art. arXiv:1312.6114, Dec. 2013.
- [12] L. McInnes, J. Healy, and J. Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv e-prints*, art. arXiv:1802.03426, Feb. 2018.
- [13] R. B. Metcalf, M. Meneghetti, C. Avestruz, F. Bellagamba, C. R. Bom, E. Bertin, R. Cabanac, F. Courbin, A. Davies, E. Decenci re, R. Flamary, R. Gavazzi, M. Geiger, P. Hartley, M. Huertas-Company, N. Jackson, C. Jacobs, E. Jullo, J. P. Kneib, L. V. E. Koopmans, F. Lanusse, C. L. Li, Q. Ma, M. Makler, N. Li, M. Lightman, C. E. Petrillo, S. Serjeant, C. Sch fer, A. Sonnenfeld, A. Tagore, C. Tortora, D. Tuccillo, M. B. Valent n, S. Velasco-Forero, G. A. Verdoes Kleijn, and G. VERNARDOS. The strong gravitational lens finding challenge. *Astronomy & Astrophysics*, 625:A119, May 2019. doi: 10.1051/0004-6361/201832797.
- [14] D. J. Schlegel, D. P. Finkbeiner, and M. Davis. Maps of Dust Infrared Emission for Use in Estimation of Reddening and Cosmic Microwave Background Radiation Foregrounds. *The Astrophysical Journal*, 500(2):525–553, June 1998. doi: 10.1086/305772.
- [15] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. CNN Features off-the-shelf: an Astounding Baseline for Recognition. *arXiv e-prints*, art. arXiv:1403.6382, Mar. 2014.
- [16] G. Stein, J. Blaum, P. Harrington, T. Medan, and Z. Lukic. Mining for strong gravitational lenses with self-supervised learning. *arXiv e-prints*, art. arXiv:2110.00023, Sept. 2021.
- [17] L. van der Maaten and G. E. Hinton. Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#)
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#)
 - (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? [\[N/A\]](#)
 - (b) Did you include complete proofs of all theoretical results? [\[N/A\]](#)
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)? [\[Yes\]](#)
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#)
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[N/A\]](#)
 - (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)? [\[Yes\]](#)
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? [\[N/A\]](#)
 - (b) Did you mention the license of the assets? [\[N/A\]](#)
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)

- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
- 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? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]