
Estimating Galactic Distances From Images Using Self-supervised Representation Learning

Md Abul Hayat

University of Arkansas
Fayetteville, AR 72701
mahayat@uark.edu

Peter Harrington

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

George Stein

Berkeley Center for Cosmological Physics
Lawrence Berkeley National Laboratory
Berkeley, CA 94720
georgestein@lbl.gov

Zarija Lukić*

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

Mustafa Mustafa*

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

Abstract

We use a contrastive self-supervised learning framework to estimate distances to galaxies from their photometric images. We incorporate data augmentations from computer vision as well as an application-specific augmentation accounting for galactic dust. We find that the resulting visual representations of galaxy images are semantically useful and allow for fast similarity searches, and can be successfully fine-tuned for the task of redshift estimation. We show that (1) pretraining on a large corpus of unlabeled data followed by fine-tuning on some labels can attain the accuracy of a fully-supervised model which requires 2-4x more labeled data, and (2) that by fine-tuning our self-supervised representations using all available data labels in the Main Galaxy Sample of the Sloan Digital Sky Survey (SDSS), we outperform the state-of-the-art supervised learning method.

1 Introduction

Sky surveys collect large amounts of data that can be used to glean information about the physical model of the universe and its evolution. One significant problem is determining 3D position of each galaxy, since its distance from the Earth is not directly measurable as are the two other spatial coordinates. Instead, one can measure the redshift of each galaxy (z) using its spectrum by looking for the departure of spectral lines from their rest-frame position (we will call this “spec- z ”). This is precise, but is a very costly procedure, as taking a spectrum requires far longer exposure times than making a photo-image of a galaxy. As a result, the vast majority of observed galaxies do not have spec- z ; instead, their redshifts are estimated from photometric data [2, 14], i.e. from images in a few passbands (think of them as the astronomical version of RGB color space, generally with ~ 5 channels). Historically, these “photo- z ” estimates commonly relied on template fitting [11] or linear regression [7] as a function of the total flux of each galaxy in different passbands, and neglected

*Corresponding authors

structural details from the images themselves. Deep neural networks however, open up new and exciting possibilities for learning photo- z 's from galaxy images using the limited number of spec- z labels. Last year, a fully-supervised convolutional neural network (CNN) was used for this task [13], showing significantly lower errors than previous models. While encouraging, this approach requires lots of labeled data.

Recent research shows that self-supervised learning techniques enable CNNs to build meaningful visual representations without needing per-sample labels for any specific downstream task [9, 4, 6]. When labels are available for only a fraction of the full dataset, the representations built via self-supervised training can then be fine-tuned for the specific vision task, and this approach is able to significantly outperform fully-supervised training [5]. Similarly, recent results in semi-supervised learning utilize a large body of unlabeled data to complement the labeled examples, and achieve state-of-the-art performance in image classification tasks as a result [17, 16].

In this work we demonstrate the advantage of self-supervised approaches for photometric redshift estimation using data from SDSS². As a first step, we establish a new fully-supervised baseline on this data. Then, using contrastive self-supervised pretraining, we show that such methods can build meaningful visual representations which are useful for a variety of tasks. By pretraining on a large corpus of unlabeled data followed by fine-tuning on labels, we can achieve the same performance as fully-supervised approaches which require 2-4x more labeled data. Finally, using the above approach and fine-tuning on the full labeled dataset, we establish a new image-based ML baseline on this dataset.

2 Methods

We use data from SDSS, a major sky survey conducted on a 2.5-meter telescope at Apache Point Observatory in New Mexico. In total, it provides photometric observations of roughly 1 billion objects in the sky, and spectra for approximately 4 million objects (stars, galaxies, and quasars).

Galaxy labels. We closely follow the process of [13] in building the labeled portion of our training dataset to enable direct comparison to their results. Using their SQL query on the `skyserver.sdss.org/CasJobs` service, we pull samples from the Main Galaxy Sample in the 12th Data Release (DR12; [1]) of the SDSS, filtering for objects classified as 'GALAXY' with dereddened petrosian magnitudes ≤ 17.8 and spectroscopic redshifts $z \leq 0.4$. For us, this query returns 547,224 objects, of which 502,977 are unique to use as labeled training examples. Executing a similar query on the 'PhotoObjAll' full photometric catalog of the SDSS, and removing duplicates which were already included in our spectroscopic sample, gives an additional body of unlabeled training examples. We imposed a cut at galactic latitude $|b| = 15^\circ$ to remove samples with high extinction near the galactic plane. We also excluded unlabeled samples with an estimated photometric redshift above 0.8 (as estimated by [3]) to eliminate objects which are very likely too distant compared to the spectroscopic sample. After imposing these cuts, our dataset of additional photometric samples contains 1,194,779 objects and is used for pre-training. The labeled training and validation dataset have 399,984 and 102,993 images respectively.

Galaxy images. SDSS photometric images contain data in 5 passbands (*ugriz*), and come background-subtracted, but are not de-reddened to account for galactic extinction. We use the Montage³ tool to query the imagery catalog in SDSS Data Release 9 (DR9), based on equatorial coordinates for each object in our dataset. We sampled a $(0.012^\circ)^2$ patch of sky centered on each object, and projected onto a 2D image with 107^2 pixels to ensure the pixel scale is as close as possible to the native pixel scale in the SDSS, 0.396 arcsec. In each image, we store the *u*, *g*, *r*, *i*, and *z* passbands as 5 color channels. Note that during training of the self-supervised model we impose random rotations and random jitter to each image before cropping out the central portion as a data augmentation, so while our images contain 107 pixels per side, the CNNs in this work only view samples of size 64^2 , which is consistent with [13]. Galaxies with smaller redshift i.e. closer to our galaxy may have a larger shape that goes beyond size 64^2 , but number of such galaxies being really small we limit the view to 64^2 for simplicity. Sample images (transformed into RGB color space) are shown in Fig. 1.

Self-supervised Learning. Recent self-supervised models like SimCLR [4, 5] and MoCo [9, 6] achieve success using a contrastive learning approach, which aims to build representations that are

² <https://www.sdss.org/> ³ <http://montage.ipac.caltech.edu/>

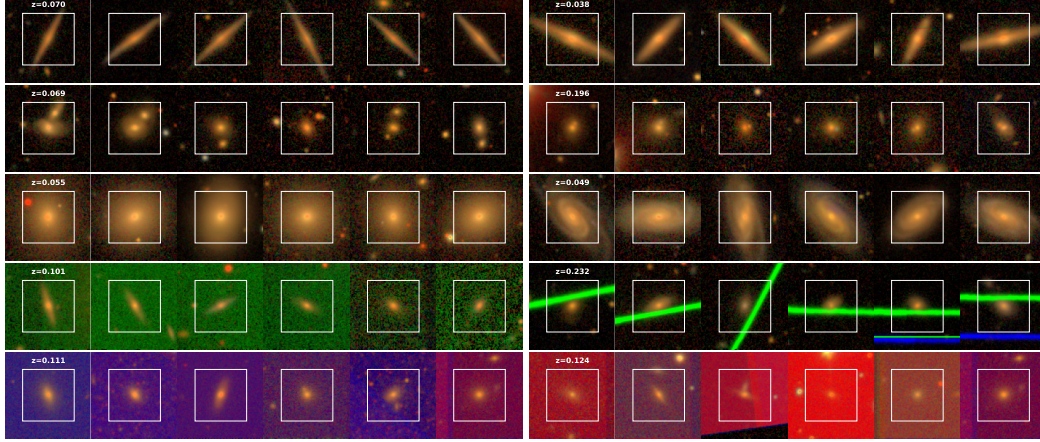


Figure 1: Reference SDSS galaxies from the validation set (leftmost panels with redshift labels) and the most similar galaxies from the training set (following 5 panels) identified through a self-supervised similarity search. White squares outline the 64^2 pixels that are “seen” by the network.

invariant under various transformations (e.g., random crops, Gaussian noise, or color jitter). These networks are trained with contrastive losses like NT-Xent as in [4] or InfoNCE in [12], which try to maximize the similarity of two augmented versions of the same image while minimizing similarity with augmented versions of other images.

In this work we use the MoCov2 framework [6], which employs randomized data augmentations to produce two different views x^q, x^k of each training image. In each training step, the differing views of each sample are passed through two independent networks, an encoder and a momentum encoder, which have identical architectures. This yields representations q and k^+ which represent a “positive” pair of samples, while “negative” examples k^- are retrieved from a queue containing previous outputs of the momentum encoder on other augmented images. The encoder network is trained via backpropagation using the InfoNCE loss

$$L_{q,k^+,\{k^-\}} = -\log \frac{\exp(q \cdot k^+ / \tau)}{\exp(q \cdot k^+ / \tau) + \sum_{k^-} \exp(q \cdot k^- / \tau)}, \quad (1)$$

where, $\tau \in (0, 1)$ is a temperature parameter. InfoNCE assigns higher similarity scores between the augmented views of the same image compared to augmented views from other images. Then, the parameters θ_k of the momentum encoder network are updated using the encoder parameters θ_q with momentum parameter m via

$$\theta_k \leftarrow m\theta_k + (1 - m)\theta_q. \quad (2)$$

The momentum update and use of a queue allow many negative examples to be seen throughout training without requiring massive batch sizes, making MoCo computationally efficient [6]. Following MoCov2, we set $m = 0.999$ and use the ResNet50 architecture [8] for the encoder and momentum encoder networks. However, we remove the first convolution and pooling layers in the ResNet and replace them with a single stride=1 convolution with 5 input feature channels to match the dimensionality of our dataset.

Data Augmentations. We create different views of training samples using the following data augmentations:

- *Galactic extinction:* To model the effects of foreground galactic dust, we introduce artificial reddening by sampling a $E(B - V)$ reddening value from $\mathcal{U}(0, 0.5)$ and applying the corresponding per-channel extinction according to the photometric calibration from [15].
- *Random rotate:* The angle of rotation is sampled from $\mathcal{U}(0, 2\pi)$.
- *Random jitter & crop:* Two integers are sampled from $\mathcal{U}(-7, 7)$ to move (jitter) the center of the image along each respective axis, then the jittered image is center-cropped to size 64^2 .
- *Gaussian noise:* We sample a scalar from $\mathcal{U}(1, 3)$ and multiply it with the aggregate median absolute deviation (MAD) of each channel (pre-computed over all training examples) to get a

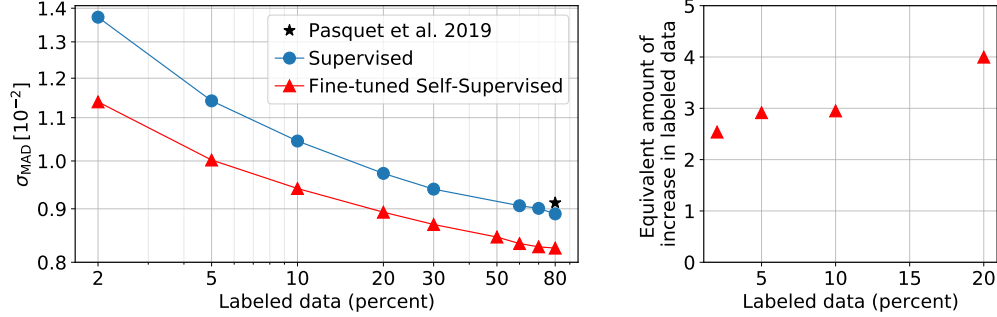


Figure 2: Our fine-tuned self-supervised approach outperforms a supervised ResNet50 at different fractions of training data (left). To achieve the same σ_{MAD} , the amount of labeled data needed for the supervised approach is 2-4x higher than for the self-supervised approach (right).

per-channel noise scale γ_c . Then, we introduce Gaussian noise sampled from $\mathcal{N}(0, \gamma_c)$ for each color channel.

3 Results

To evaluate the quality of our redshift estimates, we follow the statistics used in the literature [14, 13]. The prediction residual is defined as $\Delta z = (z_p - z_s)/(1 + z_s)$, where z_p and z_s correspond to photometric and spectroscopic redshifts respectively. From this, we compute the MAD deviation $\sigma_{\text{MAD}} = 1.4826 \times \text{MAD}(\Delta z)$ where MAD is the median of $|\Delta z - \text{median}(\Delta z)|$, and the percent η of “catastrophic” outliers with $|\Delta z| > 0.05$. To the best of our knowledge, [13] provide state-of-the-art fully-supervised photometric redshift estimates, achieving $\sigma_{\text{MAD}} = 0.00912$ and $\eta = 0.31$ using all of their labeled data. Their network is trained to be a classifier over 180 bins of photometric redshifts linearly spanning $z = (0, 0.4)$. Their estimate is computed as $z_p = \mathbb{E}(z)$ using the probability density function given by the final softmax layer, so we use the same protocol.

To properly evaluate our ResNet-based self-supervised approach, we also train fully-supervised ResNet50s on de-reddened versions of our labeled samples. This is done with 2.5%, 6.25%, 12.5%, 25%, 37.5%, 62.5%, 75%, 87.5% and 100% of the training data, with a fixed validation set. Our ResNet50, with 100% of our labeled training data, sets a new benchmark on fully-supervised estimation of photometric redshift with $\sigma_{\text{MAD}} = 0.00890$ (blue points in Fig. 2) and $\eta = 0.304$, a marginal improvement over the previous work.

Self-supervised representations. After the contrastive learning phase, galaxies are passed through the encoder network to obtain their 128 dimensional contrastive loss vectors. By computing the dot products of a query vector with those of the unlabeled training samples, and sorting by the largest values, galaxies are returned sorted by their similarity to the query. These results are depicted in Fig. 1, where the top 3 rows show examples of common galactic morphologies, while the bottom 2 have identified a number of potential observational errors that unknowingly remained after the data quality cuts were imposed. Such observational errors could be caused by foreground sources like satellites or aircraft, or by cosmic rays, leading to artifacts like the bright green lines seen in some images. It is beyond the scope of this paper, but our methodology could be used for galaxy similarity identification (see, for example, reference [10]), or for anomaly detection.

Self-supervised fine-tuning. After self-supervised pre-training on our unlabeled dataset, we partition the encoder network after the final avgpool layer (yielding a convolutional stage followed by a fully-connected layer). We fine-tune these sections on the de-reddened labeled data with a 10x smaller learning rate for the convolutional portion of the network, using the random rotations and random jitter-crop augmentations. We find that when fine-tuning our self-supervised representations on the fractions of the training dataset as explained above, we outperform the fully-supervised ResNet50 baseline for all fractions, as shown in Fig. 2. Our self-supervised network, fine-tuned on 100% of the labeled training data, outperforms both our fully-supervised baseline and the previous state-of-the-art result with $\sigma_{\text{MAD}} = 0.00825$ and $\eta = 0.209$. Importantly, for the intermediate fractions of the training dataset we find that our self-supervised network achieves the equivalent accuracy of a

fully-supervised network while using 2-4x less labeled data (right panel in Fig. 2). We have also confirmed that the prediction bias $\langle \Delta z \rangle$, defined as the mean of the residuals, is negligibly small in our work ($\langle \Delta z \rangle < 10^{-4}$) as it was in [13].

4 Conclusion

In this work we have presented a first effort in leveraging self-supervised learning to process unlabeled photometric data, showing that the resulting visual representations are semantically useful and can be fine-tuned to surpass the performance of fully-supervised photometric redshift estimation models. We demonstrate that our model performs well even in the case of limited data labels, which is extremely valuable given the relative cost of acquiring spectroscopic measurements for a large sample of galaxies. Upcoming large-scale sky surveys from the ground and space, like the Vera C. Rubin Observatory⁴ and Euclid⁵, are projected to image tens of billions of galaxies over the next decade. With those surveys in mind, our results show great promise for self-supervised learning methods to assist in deriving more precise photometric redshift estimates, helping address fundamental physics and cosmology questions on the nature and properties of dark energy, dark matter and gravity.

Broader Impact

As a component of this publication we intend to (pursuant to SDSS approval) publicly release our processed version of galaxy images and redshift labels, as well as all networks used in this research. Data we have pre-processed is immediately ready for different learning tasks and easy to manipulate with e.g. python codes, which is not the case with the more complex SDSS data server. Our full code release makes it straightforward to reproduce our results and build upon it. This reduces the entry threshold for future machine learning explorations with this interesting data set, especially for researchers with no background in astronomy, and for students who are not in top-tier research universities.

References

- [1] S. Alam, F. D. Albareti, C. A. Prieto, F. Anders, S. F. Anderson, T. Anderton, B. H. Andrews, E. Armengaud, É. Aubourg, S. Bailey, et al. The eleventh and twelfth data releases of the sloan digital sky survey: final data from sdss-iii. *The Astrophysical Journal Supplement Series*, 219 (1):12, 2015.
- [2] W. A. Baum. Photoelectric Magnitudes and Red-Shifts. In G. C. McVittie, editor, *Problems of Extra-Galactic Research*, volume 15 of *IAU Symposium*, page 390, Jan. 1962.
- [3] R. Beck, L. Dobos, T. Budavári, A. S. Szalay, and I. Csabai. Photometric redshifts for the SDSS Data Release 12. *Monthly Notices of the Royal Astronomical Society*, 460(2):1371–1381, Aug. 2016. doi: 10.1093/mnras/stw1009.
- [4] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [5] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020.
- [6] X. Chen, H. Fan, R. Girshick, and K. He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [7] A. J. Connolly, I. Csabai, A. S. Szalay, D. C. Koo, R. G. Kron, and J. A. Munn. Slicing Through Multicolor Space: Galaxy Redshifts from Broadband Photometry. *Astronomical Journal*, 110: 2655, Dec. 1995. doi: 10.1086/117720.
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. *arXiv e-prints*, art. arXiv:1512.03385, Dec. 2015.

⁴ <https://www.lsst.org/> ⁵ <https://www.euclid-ec.org/>

- [9] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.
- [10] A. Hocking, J. E. Geach, Y. Sun, and N. Davey. An automatic taxonomy of galaxy morphology using unsupervised machine learning. *Monthly Notices of the Royal Astronomical Society*, 473(1):1108–1129, Jan. 2018. doi: 10.1093/mnras/stx2351.
- [11] E. D. Loh and E. J. Spillar. Photometric Redshifts of Galaxies. *The Astrophysical Journal*, 303:154, Apr. 1986. doi: 10.1086/164062.
- [12] A. v. d. Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [13] J. Pasquet, E. Bertin, M. Treyer, S. Arnouts, and D. Fouchez. Photometric redshifts from sdss images using a convolutional neural network. *Astronomy & Astrophysics*, 621:A26, 2019.
- [14] M. Salvato, O. Ilbert, and B. Hoyle. The many flavours of photometric redshifts. *Nature Astronomy*, 3(3):212–222, 2019.
- [15] E. F. Schlafly and D. P. Finkbeiner. Measuring reddening with sloan digital sky survey stellar spectra and recalibrating sfid. *The Astrophysical Journal*, 737(2):103, 2011.
- [16] H. Touvron, A. Vedaldi, M. Douze, and H. Jégou. Fixing the train-test resolution discrepancy: FixEfficientNet. *arXiv e-prints*, art. arXiv:2003.08237, Mar. 2020.
- [17] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le. Self-training with Noisy Student improves ImageNet classification. *arXiv e-prints*, art. arXiv:1911.04252, Nov. 2019.