Unsupervised Learning of Deep Features Through Best-Fits for Observational Cosmology

Christopher Murray
École Polytechnique Paris
Palaiseau, Essonne, France
christophe.murray@polytechnique.edu

Sébastien Fabbro
NRC Herzberg
Victoria, BC, Canada
sfabbro@uvic.ca

Kwang Moo Yi
University of Victoria
Victoria, BC, Canada
kyi@uvic.ca

Abstract

We propose a novel method based on Generative Adversarial Networks (GAN) and optimization-based inference to learn deep features for observational cosmology. Measuring the universe's properties often involve distances and redshifts, which are very costly measurements. This scarcity has lead to machine learning pipelines relying mostly on handcrafted features. In this work, we propose an alternative that takes advantage of the vast amount of unlabeled data that is already available. Specifically, we learn a latent space of galaxies through a GAN for sharp image reconstruction, which we further constrain with a latent space reconstruction. We demonstrate qualitatively that our method successfully learns to recreate astronomical sources. We further show quantitatively that the deep features learned through our method are able to deliver promising results in estimating redshifts from galaxy multi-channel images.

1 Introduction

Physical models of cosmology imply tight relations between the geometry and kinematics of the universe, often quantified from the measure of distances and redshifts using wide field imaging and spectroscopy survey telescope data. Given the difficulty of measuring extragalactic distances, only a tiny fraction of the billions of detected astronomical sources in recent public surveys such as the Sloan Digital Sky Survey [3] or the Dark Energy Survey [1]. We can attain $\approx$6% precision using very scarce type Ia supernova events to estimate distances, but even after 20 years since their use to reveal we live in an accelerating universe, we barely have 1,000 reliable distances [22]. Precise measurements of redshifts are more numerous but still constitute only a tiny fraction of the available imaged galaxies, and often handcrafted features are used to measure the rest of the data with a trade-off on precision [20], and at higher redshifts, physical based models often perform best.

Very recently, attempts were made to change this situation [7, 17] with Convolutional Neural Networks (CNNs) showing promising results, the scope of the problem was limited to nearby ($z < 0.8$) sources, which allows for a larger amount of labeled data (100k) and a fully supervised approach.

In this work, we explore whether one could learn data-driven features for galaxies without much restriction, and apply them to understand galaxies with few labeled samples – e.g. redshift for galaxies that are distant ($z > 2$). To do so, we propose a method based on Generative Adversarial Networks (GAN) [11] and learn an embedding space that could be used as deep features. While other groups have extended GANs e.g. [18, 10] to galaxy image generation to calibrate cosmological surveys, efforts are under way to explore how to infer physical properties from the generative models. Here we focus how we could reach larger range of redshifts by integrating very deep multi-survey data sets.

In more detail, we train a modified Least Squares GAN (LSGAN) [14] to learn an embedding space of galaxies from *unlabeled* galaxy images available from the Hyper Suprime Camera Data Release

Second Workshop on Machine Learning and the Physical Sciences (NeurIPS 2019), Vancouver, Canada.

DR2 [2]. To extract deep features using this learned embedding space, we train a network that learns to re-embed generated images to the original representation similar to [9, 4]. However, unlike [9, 4], we simply train the *embedder* to reproduce the original noise signal that was used to create a synthetic image, instead of trying to embed real galaxy images. By doing so, we learn a fully unsupervised deep feature representation for the observed astronomical field.

The novelty of our method lies in the inference stage when we extract deep features. We optimize to infer the latent embeddings for a given input image. In more details, with the initial embedding given by the *embedder*, we iteratively update the latent space embeddings to find one that minimize the reconstruction error the given image. As we show later in Fig. 1, this allows us to obtain accurate reconstructions, whereas a simple feed-forward approach with the embedder is not able to do so.

We further empirically show that our deep features can be used to regress the redshifts for distant galaxies. With our deep features, we use the TPOT package [16], which automatically derives a regression pipeline based on Gradient Boosting [5]. This provides comparable results to the state of the art [21]. Given that our deep features are never explicitly related to any physical characteristics of the astronomical sources during training, this demonstrates the potential in learning the hidden physical properties of galaxies by simply observing much data.

2 Proposed Method

Our network is composed of a typical GAN framework, with an additional embedder that tries to recover the embedding space from the generated fake image. We first discuss our network architecture, then the training objective, and finally our novel method that utilizes the trained generator and the embedder to retrieve the deep features.

2.1 Architectures

Our design choice are empirically motivated from our attempts to form a stable GAN framework.

Generator. We draw a 64D space $\mathbf{z}$ from a normal distribution, which is then fed to a fully connected layer with 4098 output neurons. We then transform the output into a 2×2 image with 1024 channels. We further apply 3 residual blocks [12] to expand this image into a 16×16 image with 128 channels, where each residual block is composed of a transposed convolution with a stride of two for upsampling followed by ReLU activation. The pass-through route is composed of a simple transposed convolution with no activation. The number of channels is halved at each upsampling. We found empirically that it was beneficial to *not* employ any normalization to the generator. We finally apply a simple convolution with no activation to turn the image into a 32×32 image with 5 channels, which is what our observed data looks like.

Discriminator and Embedder. We apply an architecture that mirrors the generator – convolutional layers in place of transposed convolutions and in reverse order as the generator – but with LeakyReLU activations. For all layers we apply Spectral Normalization [15]. In addition, we apply Group Normalization [24] for each convolution layer with four groups. The difference between the discriminator and the embedder is in the fully-connected layer, which the embedder has 64 output neurons and the discriminator has one output neuron, respectively.

2.2 Training

To form a GAN framework and learn an embedding space, we apply the LSGAN objective[1] but with an additional objective to train the embedder. The additional objective that we apply is to reconstruct the embeddings with an embedder. If we denote the latent space embedding as $\mathbf{z}$, the generator with parameters $\mathbf{\Phi}$ as $G_{\mathbf{\Phi}}$, and the embedder as $E_{\mathbf{\Theta}}$, where $\mathbf{\Theta}$ is the parameters of this embedder, we define an additional loss term $\mathcal{L}_E$ for training as

$$\mathcal{L}_E = \mathbb{E}_{\mathbf{z} \sim \mathcal{N}} \left[\left\| \mathbf{z} - E_{\mathbf{\Theta}} \left(G_{\mathbf{\Phi}} \left(\mathbf{z} \right) \right) \right\|_2^2 \right], \tag{1}$$

where $\mathcal{N}$ is a normal distribution. Note that in [8], the embedder was trained to operate on real images, and create a pair of latent embedding and image, which was then distinguished by a discriminator.

[1] Because of spatial constraints, we omit this objective and refer the interested readers to [14].

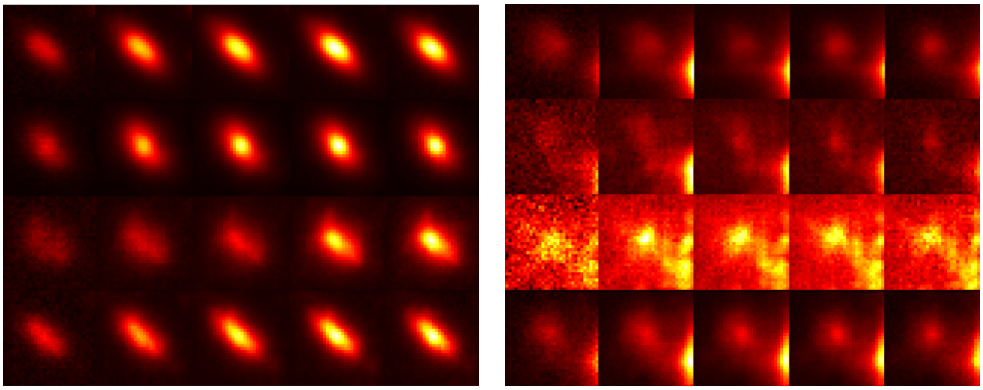

Figure 1: **Qualitative examples.** All five channels are shown as columns for the two galaxy images. (First row) actual observations, (second row) reconstructions using only the embedder, (third row) reconstructions using only the optimization starting from a random embedding, (fourth row) our method. Our method is able to provide accurate reconstructions. As shown, optimization alone fails in both cases and the embedder alone fails for the left example.

This creates difficulty in training, as different data types are fed to the discriminator, and the user further needs to balance between the embedding and the image for training the discriminator. Our approach alleviates this complexity and keeps things simple. This further allows us to even train the embedder separately if needed.

However, as shown in Fig. 1, an embedder trained in this way is not able to deliver good embeddings, which is why we opt for a unique inference flow that we detail in the next subsection.

2.3 Inference

The uniqueness of our method lies in the inference phase. Our method starts from the intuition that our ultimate goal with the embedder is to be able to find the latent space that best reconstructs a given image example. Therefore, instead of aiming to train a perfect embedder that may not be possible, starting from the initial estimate provided by the embedder, we explicitly minimize this reconstruction error, with respect to the latent embeddings. Mathematically, if we denote the image we wish to obtain the latent embeddings for as $\mathbf{I}$, we obtain our latent embedding $\hat{\mathbf{z}}$ as

$$\hat{\mathbf{z}} = \underset{\mathbf{z}}{\operatorname{argmin}} \left\| \mathbf{I} - G_{\boldsymbol{\Theta}}\left(\mathbf{z}\right) \right\|_2^2 . \tag{2}$$

To solve this equation and retrieve $\hat{\mathbf{z}}$, we use the Adam optimizer [13] with a learning rate of 0.25.

Notice that there is no embedder in Eq. (2). This does not imply that an embedder is unnecessary. Rather, the optimization problem defined in Eq. (2) is highly non-convex, while the optimization we perform is based on gradient descent. It is therefore important that the embedder provides a proper initial guess to the location; see Fig. 1. Still, this requirement is not as strict as the alternative where one aims to learn a perfect embedder.

3 Results

3.1 Data preparation

Our data sets consist of very deep extragalactic fields observed in multiple filters u,g,r,i,z, with the Subaru and the Canada France Hawaii Telescopes. Hundreds of exposures we registered and stacked for each filter [19, 21] to produce higher larger signal to noise images, and from the catalogue of detected sources on the stacked images, we extract 200k 32×32 cutouts centered on a source, which happens to be most of the time a resolved galaxy but on rarer occasion a star or quasar. We do not apply any selection criterion on the nature, luminosity or shape of the sources.

For training our modified LSGAN, images with missing or corrupted information are discarded, and the remaining data is normalized, per channel, by the average mean and standard deviation of all input galaxy field images.

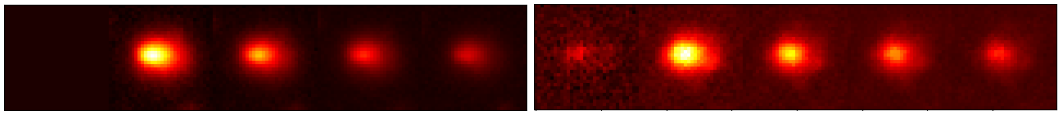

Figure 2: **Example of impainting the u-channel.** The u-channel of the observation (left) for some reason has been destroyed in the dataset. Even so, our method is able to impaint the u-channel (right) by finding the latent embedding that best reconstructs the other channels *without any* re-training.

3.2 Reconstructing field images

We provide qualitative examples for reconstructing field images in Fig. 1. Comparing results with and without the embedder reveals the importance of a good-enough initial guess, empirically demonstrating that the solution space for Eq. (2) is indeed non-convex. The poor results without the embedding optimization show that an embedder alone is unable to provide an informative embedding.

Besides simple reconstructions, we provide another example in Fig. 2, which demonstrates the potential of the proposed method. In Fig. 2, we show that *without any re-traning* our method is able to find the latent embedding that best fits the other four channels, recreating the u-channel image in doing so. This highlights the potential of our method being able to adapt to situations where only partial measurements are given.

3.3 Predicting galaxy redshifts from field images

To demonstrate the usefulness of our deep features, we use an off-the-shelf automatic regressor [16] with our deep features. The redshifts measured through spectroscopic measurements are taken as the ground-truth labels from the COSMOS [23] and CLAUDS [21] surveys. We evaluate the quality of predictions quantitatively with four standard metrics for cosmology: R^2; Bias, or $\langle \Delta z \rangle$, which is a average of Δz following the definition by [6]; Mean Absolute Deviation, σ_{MAD}; fraction η of outliers, with outliers defined as having $|\Delta z| > 5\sigma_{MAD}$ as in [21]. We compare our method against a state-of-the-art method by Sawicki *et al.* [21][2], and a CNN regressor using the 50 layer architecture proposed in [25], using the same field images as input. Table 1 summarizes these results. Our method provides outperforms Sawicki *et al.* [21] in terms of R^2 and η, but performs worse in terms of $\langle \Delta z \rangle$ and σ_{MAD}. Note that using a wide ResNet completely fails as shown by the R^2 and σ_{MAD} – this also makes the η value invalid as none of samples are considered outliers. While there is room for improvement, given that our deep features were trained fully unsupervised, it demonstrates the potential of our deep features.

Table 1: **Quantitative results.** We show the best method in bold and underline the second best. See text for details.

Method	R^2	$< z >$	σ_{MAD}	η
Wide ResNet [25]	0.0110	-0.0277	0.253	**0.0**
Sawicki *et al.* [21]	<u>0.608</u>	**0.00921**	**0.0589**	0.450
Our Method	**0.633**	<u>-0.0235</u>	<u>0.130</u>	<u>0.389</u>

4 Conclusions and Future Work

We have proposed a method based on GANs and best-fits to learn deep features for cosmology form field images. We demonstrated that the proposed method is able to obtain embeddings that could be used to reconstruct field images with deep networks. We further showed an early demonstration of using our deep feature to regress redshifts.

While we have only demonstrated it for redshift of galaxies, we expect our deep features to carry much more than just redshifts, such as distances or image features with physical interest. While the derived deep features are fully unsupervised and data driven, we plan to further incorporate semi-supervised approaches to enhance these deep features even further.

Acknowledgements

This work was partially supported by the NSERC Discovery Grant (RGPIN-2018-03788, DGECR-2018-00426), and by systems supplied by Compute Canada.

[2]We thank the authors for providing their redshift estimates.

References

[1] T. Abbott et al. The Dark Energy Survey: Data Release 1. *The Astrophysical Journal Supplement Series*, 239(2):18, Dec 2018.

[2] H. Aihara et al. Second Data Release of the Hyper Suprime-Cam Subaru Strategic Program. *arXiv Preprint*, 2019.

[3] S. Alam 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, Jul 2015.

[4] M. I. Belghazi, S. Rajeswar, O. Mastropietro, N. Rostamzadeh, J. Mitrovic, and A. C. Courville. Hierarchical adversarially learned inference. *arXiv Preprint*, 2018.

[5] T. Chen and C. Guestrin. XGBoost: A scalable tree boosting system. In *Conference on Knowledge Discovery and Data Mining*, 2016.

[6] J. G. Cohen, D. W. Hogg, R. Blandford, L. L. Cowie, E. Hu, A. Songaila, P. Shopbell, and K. Richberg. Caltech faint galaxy redshift survey. x. a redshift survey in the region of the hubble deep field north. *The Astrophysical Journal*, 538(1):29–52, jul 2000.

[7] A. D'Isanto and K. L. Polsterer. Photometric redshift estimation via deep learning. Generalized and pre-classification-less, image based, fully probabilistic redshifts. *Astronomy & Astrophysics*, 609:A111, Jan 2018.

[8] J. Donahue, P. Krähenbühl, and T. Darrell. Adversarial Feature Learning. In *International Conference on Learning Representations*, 2017.

[9] J. Donahue and K. Simonyan. Large Scale Adversarial Representation Learning. *arXiv Preprint*, 2019.

[10] L. Fussell and B. Moews. Forging new worlds: high-resolution synthetic galaxies with chained generative adversarial networks. *Monthly Notices of the Royal Astronomical Society*, 485(3):3203–3214, May 2019.

[11] I. Goodfellow, J. Pouget-abadie, M. Mirza, B. Xu, D. Warde-farley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.

[12] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.

[13] D. Kingma and J. Ba. Adam: A Method for Stochastic Optimisation. In *International Conference on Learning Representations*, 2015.

[14] X. Mao, Q. Li, H. Xie, R. Lau, Z. Wang, and S. Smolley. Least Squares Generative Adversarial Networks. In *International Conference on Computer Vision*, pages 2794–2804, 2017.

[15] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.

[16] R. S. Olson, R. J. Urbanowicz, P. C. Andrews, N. A. Lavender, L. C. Kidd, and J. H. Moore. *Applications of Evolutionary Computation: 19th European Conference, EvoApplications 2016, Porto, Portugal, March 30 – April 1, 2016, Proceedings, Part I*, chapter Automating Biomedical Data Science Through Tree-Based Pipeline Optimization, pages 123–137. Springer International Publishing, 2016.

[17] 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, Jan 2019.

[18] S. Ravanbakhsh, F. Lanusse, R. Mandelbaum, J. Schneider, and B. Poczos. Enabling Dark Energy Science with Deep Generative Models of Galaxy Images. *arXiv e-prints*, page arXiv:1609.05796, Sep 2016.

[19] M. Roman et al. Dependence of Type Ia supernova luminosities on their local environment. *Astronomy and Astrophysics*, 615:A68, Jul 2018.

[20] M. Salvato, O. Ilbert, and B. Hoyle. The many flavours of photometric redshifts. *Nature Astronomy*, 3:212–222, Jun 2019.

[21] M. Sawicki et al. The CFHT Large Area U-band Deep Survey (CLAUDS). *arXiv Preprint*, 2019.

[22] D. Scolnic et al. The Complete Light-curve Sample of Spectroscopically Confirmed SNe Ia from Pan-STARRS1 and Cosmological Constraints from the Combined Pantheon Sample. *Astrophysical Journal*, 859(2):101, Jun 2018.

[23] N. Scoville et al. The Cosmic Evolution Survey (COSMOS): Overview. *The Astrophysical Journal Supplement*, 172(1):1–8, Sep 2007.

[24] Y. Wu and K. He. Group Normalization. In *European Conference on Computer Vision*, 2018.

[25] S. Zagoruyko and N. Komodakis. Wide residual networks. In *British Machine Vision Conference*, 2016.