
Solving Inverse Problems for Spectral Energy Distributions with Deep Generative Networks

Agapi Rissaki

National Technical University of Athens
arissaki@corelab.ntua.gr

Orestis Pavlou

European University Cyprus
o.pavlou@euc.ac.cy

Dimitris Fotakis

National Technical University of Athens
fotakis@cs.ntua.gr

Vicky Papadopoulou

European University Cyprus
v.papadopoulou@euc.ac.cy

Andreas Efstathiou

European University Cyprus
a.efstathiou@euc.ac.cy

Abstract

We propose an end-to-end approach for solving inverse problems for a class of complex astronomical signals, namely Spectral Energy Distributions (SEDs). Our goal is to reconstruct such signals from scarce and/or unreliable measurements. We achieve that by leveraging a learned structural prior in the form of a Deep Generative Network. Similar methods have been tested almost exclusively for images which display useful properties (e.g., locality, periodicity) that are implicitly exploited. However, SEDs lack such properties which make the problem more challenging. We manage to successfully extend the methods to SEDs using a Generative Latent Optimization model trained with significantly fewer and corrupted data.

1 Introduction

In astrophysics, distributions constructed by energy measurements in different wavelengths, namely Spectral Energy Distributions (SEDs), are important tools for studying the physical properties and evolution of astronomical objects. SEDs can be used for example to determine the luminosity of astronomical objects, the rate at which galaxies form new stars or the rate at which supermassive black holes accrete mass to generate energy in quasars [1, 2, 3]. However, the measurement process is prone to statistical (random) as well as systematic errors, such as background and foreground noise interference, *i.e.*, atmospheric absorption and distortion, opaque/obscuring dust, *etc.* Due to these factors, as well as technical limitations, such as camera sensor sensitivity, cooling, resolution, *etc.*, SEDs are collected in scarce, often incomplete datasets. SEDs are compared to physical models in order to find the best-fit model(s), which provides an insight into the underlying physical processes and properties of the target. This highlights the importance of expanding the range and improving the accuracy of the available data points. In the literature, computational methods have been widely used to enhance SEDs and handle the experimental error [4].

In recent years, deep learning has proven to be an important tool for enhancement of real data and in general for solving inverse problems, where the goal is to reconstruct or correct a signal given an incomplete and/or noisy version. Specifically for astronomical data, deep learning techniques have been used mainly for astronomical imaging, such as deblending images of galaxies [5] or image enhancement [6]. For SEDs, deep learning has been used in forward problems such as feature

extraction [7], but not inverse problems. In this paper we use well-known deep learning techniques adjusted appropriately in order to solve various inverse problems for SEDs.

The method we apply is data-driven and utilizes Deep Generative Models as learned structural priors. More specifically, models like Variational AutoEncoders (VAEs) [8] and Generative Adversarial Networks (GANs) [9], trained on large datasets (most frequently of images) are able to extract information about the underlying data distribution and generate realistic samples. These models, once trained, can be used as structural priors for solving inverse problems [10]. Thus, this methods requires us to train a high-quality generative network which can model realistic SEDs, with properties such as high-frequency, irregularity *etc.* In this paper, we use the Generative Latent Optimization framework (GLO) [11] to train a deep generative network suitable for our needs. The framework allows us to train a high-quality generative network with more flexibility than a VAE and at the same time offers training efficiency unlike GANs, which are notoriously hard to train.

In order to train a generative network any state-of-the-art method requires a high-quality large dataset. However, for the case of SEDs these prerequisites are unrealistic since the measurement procedure contains innate error, incompleteness and is particularly expensive. To overcome the issue of erroneous and/or incomplete samples we propose an end-to-end approach: (1) a preprocessing step where we utilize classical computational methods for enhancement, *e.g.*, iterative PCA [4], (2) the deep learning method described above. Our approach is useful for a variety of inverse problems and it can mitigate the long-term cost of solving such problems for SEDs. Furthermore, it is expected to improve overall performance on these problems even with significant corruption and/or incompleteness by leveraging the powerful generalization property as well as the robustness of a deep generative network.

2 Method

Suppose, we collect measurements of the form:

$$y = Ax^* + \eta, \quad (1)$$

where A is a measurement matrix and η a noise vector. Our goal is to reconstruct the signal x^* , given A and η , thus solve the linear inverse problem. This formulation usually refers to compressed sensing (where we assume few measurements taken) but can be also used to model several real-world problems concerning SEDs. We tackle the problem for the case of SEDs using a deep generative network as a structural prior [10], a method that has been successfully applied to natural images.

2.1 Building the Generative Network

To build our generative network we use the Generative Latent Optimization (GLO) framework [11], which allows us to train a relatively large generator (sufficiently over-parametrized) in order to achieve good generalization [12]. The framework is based on the manipulation of the generator’s latent space as well as its parameters using a simple reconstruction loss. We use the GLO framework as an alternative to GANs which are trained via an adversarial optimization scheme. Unlike GLO, which consists of a simple loss minimization back-propagated to the latent space, GANs should ideally converge to an (approximate) equilibrium which is not guaranteed and/or requires excessive resources [13]. Thus, when training GANs in practice it is common to examine the generated samples and stop the training when they are satisfactory. In the case of images this technique can be easily applied, but for SEDs this is not feasible. In fact, we use the ability to solve inverse problems as a proxy to evaluate our trained generator.

Let us examine the training procedure more closely. We train the generator $G : \mathcal{Z} \rightarrow \mathcal{X}$, where \mathcal{Z} denotes the latent space and \mathcal{X} the underlying class of SEDs which is described by the training set $\{x_i\}_{i=1}^N$. Prior to training, we randomly initialize the latent codes $z_i \in \mathcal{Z}$ from a multi-dimensional Normal distribution and pair them with each of the samples x_i . During training, the generator’s parameters and the latent codes $\{z_i\}_{i=1}^N$ are jointly optimized, as described by (2). The optimization is driven by a simple reconstruction loss $\mathcal{L}(\cdot)$, which in our case is Mean Squared Error (MSE).

$$\min_G \frac{1}{N} \sum_{i=1}^N \min_{z_i \in \mathcal{Z}} [\mathcal{L}(G(z_i), x_i)] \quad (2)$$

More specifically, the gradient of the loss function with respect to the parameters of the generator and the latent code is back-propagated all the way through the network and to the latent space. This training procedure makes the latent space more structurally meaningful and suitable for reconstruction. To promote this feature, we project the latent codes onto the unit sphere during training [11].

2.2 Reconstruction

Given the generative network $G(\cdot)$, the estimated solution of an inverse problem (1) could be $\hat{x} = G(\hat{z})$ where:

$$\hat{z} = \arg \min_{z \in \mathcal{Z}} \frac{1}{m} \|AG(z) - y\|_2^2 \quad (3)$$

In other words, we (approximately) optimize the latent code \hat{z} such that the corresponding signal \hat{x} matches the measurements y . We optimize \hat{z} by back-propagating the gradient of the reconstruction loss through $G(\cdot)$ [10]. Note that we have to project z onto the unit sphere, similarly to training. In a different approach [10], instead of explicitly projecting z onto the unit sphere, we can apply a regularization to implicitly restrict z as follows:

$$\min_{z \in \mathcal{Z}} \frac{1}{m} \|AG(z) - y\|_2^2 + R(z), \quad (4)$$

where $R(z) = \lambda \|z\|_2^2$ is the regularization term and λ a balance hyperparameter.

3 Experiments

3.1 Data

We apply our approach to Sloan Digital Sky Survey (SDSS) spectra ¹. Specifically, we use the preprocessed SDSS Corrected Spectra dataset offered by the astroML library [14]. The dataset contains SEDs for 4000 galaxies moved to restframe, preprocessed with iterative PCA and resampled to 1000 wavelengths (3000 – 8000Å). Although the preprocessing is imperfect, leading to outlier values, our deep learning approach still displays great performance due to its robustness. Notice that the original SDSS dataset consists of innately incomplete and/or corrupted SEDs, due to the nature of the measurement process. Thus, the original SEDs cannot be used directly for evaluation purposes because we would lack the ground truth. Instead, we consider part of the corrected SEDs produced by the preprocessing step as test data (10% of the preprocessed dataset), which we use for comparisons in Section 3.3.

3.2 Training

We train a Feed-forward neural network with 7 hidden layers and leakyReLU activations (except for the output layer). We use 90% of the preprocessed dataset as our training data and train our network for 10000 epochs with batches of 64 spectra. We use Adam optimization [15] with learning rate 0.1 for the network’s parameters and 0.01 for the latent codes as well as 1d-batch normalization to accelerate the training procedure [16]. We choose a simple Mean Squared Error (MSE) as our loss function and we also apply weight decay to avoid overfitting.

Our spectra consist of measurements for 1000 wavelengths. We choose 50 dimensions for the latent space, which are sufficient for the representation and allow for efficient training and reconstruction. For the reconstruction, we limit the optimization procedure to 1000 epochs and choose a configuration similar to training. The project is developed using PyTorch [17].

3.3 Results

We evaluate our approach, both qualitatively and quantitatively, for different inverse problems, by artificially injecting realistic corruption and/or incompleteness to our test data. For the qualitative evaluation (Figure 1), we examine the performance of our algorithm on inverse problems with missing information. More specifically, the missing information corresponds to either a continuous window of

¹<https://www.sdss.org/dr12/spectro/>

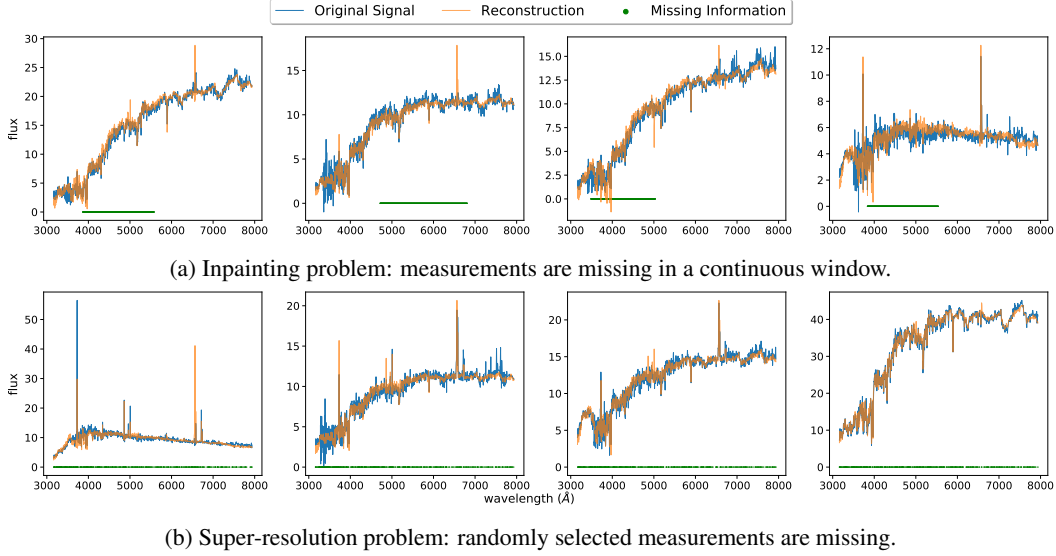


Figure 1: The original SED signals and their reconstruction for 40% missing information in inpainting (top) and super-resolution (bottom) settings.

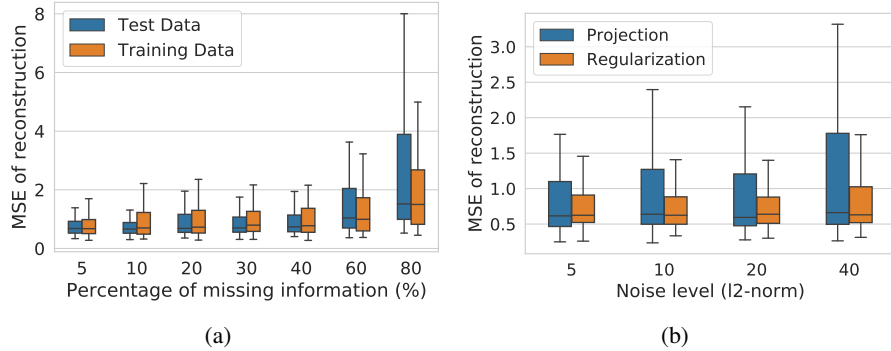


Figure 2: MSE of reconstruction for 100 randomly selected signals on: (a) inpainting, for different levels of missing information (%). (b) denoising, for different levels of added noise (l_2 -norm).

missing values (inpainting) or randomly chosen values throughout the entire signal (super-resolution). For each problem we randomly select four SED signals from our test set, then apply the appropriate masking and produce a reconstruction. The missing information in both cases is chosen to be 40% of the total number of measurements that compose each SED. We can see that for both problems, the reconstruction closely follows the trajectory of the original signal and in most cases predicts the high-frequency changes and large spikes.

In Figure 2, we examine quantitatively the performance of our algorithm on the problems of (a) inpainting and (b) denoising (that is removing added noise drawn by a normal distribution). For each configuration, we show the reconstruction MSE of 100 randomly selected SEDs (excluding measurements that fall outside 1.5 times the interquartile range). In Figure 2(a), we examine the performance for different levels of missing information and compare between SEDs drawn from test and training data. In both cases the MSE of the vast majority of signals is particularly low for up to 40% of missing information. For reasonable percentages of missing information the performance on test data is on par with the training data. For a sufficient proportion of the examined signals, this trend persists even for larger percentages (see median values). Given that our generative network was optimized to represent the training data, this shows a considerable generalization capability, which is crucial for the effectiveness of our approach. In Figure 2(b), we examine the performance for different levels of added noise and compare between our reconstruction methods, the explicit projection (eq. 3) and the regularization (eq. 4). We can see that for all levels of added noise the MSE is particularly

low, which indicates notable performance on denoising. Furthermore, when regularization is utilized we observe better error concentration, which can be attributed to the flexibility it offers to the reconstruction process.

4 Conclusion and Future Work

We presented an end-to-end deep learning solution for various inverse problems concerning Spectral Energy Distributions (SEDs). Our approach relies on a deep generative network, tailored to the particular properties of SEDs, as a structural prior leveraging its generalization capability. Our preliminary results show promising performance on realistic inverse problems. We are working to extend this project to diverse and more demanding SED families *e.g.*, for different parts of the spectrum. Another future direction involves transfer learning techniques, as well as ensemble learning in order to extend our approach to data that are even more incomplete. Finally, we could augment our method using bi-directional training in order to simultaneously extract information regarding the astrophysical objects we study. This idea draws from recent research on invertible neural networks for inverse problems [18].

Broader Impact

This project will have broad impact in the effort to interpret the SEDs which will be made available with a number of current and future ground-based and space missions such as LSST, Euclid, JWST and SPICA. Although the examples used in this work concentrate on the optical part of the spectrum, the same method can also be used on SEDs which cover the whole spectrum of galaxies from the ultraviolet to the radio. Such studies of the complete SEDs of galaxies are now recognized as essential for a complete understanding of the processes that control galaxy formation and evolution (e.g. [3, 19]).

References

- [1] Andreas Efsthathiou and Michael Rowan-Robinson. Dusty discs in active galactic nuclei. *Monthly Notices of the Royal Astronomical Society*, 273(3):649–661, 1995.
- [2] Andreas Efsthathiou, Michael Rowan-Robinson, and Ralf Siebenmorgen. Massive star formation in galaxies: radiative transfer models of the uv to millimetre emission of starburst galaxies. *Monthly Notices of the Royal Astronomical Society*, 313(4):734–744, 2000.
- [3] Michael Rowan-Robinson and et al. Spectral energy distributions and luminosities of galaxies and active galactic nuclei in the spitzer wide-area infrared extragalactic (swire) legacy survey. *The Astronomical Journal*, 129:1183–1197, March 2005.
- [4] Jakob Walcher, Brent Groves, Tamás Budavári, and Daniel Dale. Fitting the integrated spectral energy distributions of galaxies. *Astrophysics and Space Science*, 331(1):1–51, Aug 2010.
- [5] Alexandre Boucaud, Marc Huertas-Company, Caroline Heneka, Emille E O Ishida, Nima Sedaghat, Rafael S de Souza, Ben Moews, Hervé Dole, Marco Castellano, Emiliano Merlin, and et al. Photometry of high-redshift blended galaxies using deep learning. *Monthly Notices of the Royal Astronomical Society*, 491(2):2481–2495, Dec 2019.
- [6] Francois Lanusse, Peter Melchior, and Fred Moolekamp. Hybrid physical-deep learning model for astronomical inverse problems, 2019.
- [7] Frontera-Pons, J., Sureau, F., Bobin, J., and Le Floch, E. Unsupervised feature-learning for galaxy sed with denoising autoencoders. *A&A*, 603:A60, 2017.
- [8] Diederik P. Kingma and Max Welling. An introduction to variational autoencoders. *CoRR*, abs/1906.02691, 2019.
- [9] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.

- [10] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G. Dimakis. Compressed sensing using generative models. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, pages 537–546. JMLR.org, 2017.
- [11] Piotr Bojanowski, Armand Joulin, David Lopez-Paz, and Arthur Szlam. Optimizing the latent space of generative networks. 2017.
- [12] Behnam Neyshabur, Srinadh Bhojanapalli, David Mcallester, and Nati Srebro. Exploring generalization in deep learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5947–5956. Curran Associates, Inc., 2017.
- [13] Frans A. Oliehoek, Rahul Savani, Jose Gallego-Posada, Elise van der Pol, and Roderich Groß. Beyond local nash equilibria for adversarial networks. *CoRR*, abs/1806.07268, 2018.
- [14] J.T. Vanderplas, A.J. Connolly, Ž. Ivezić, and A. Gray. Introduction to astroml: Machine learning for astrophysics. In *Conference on Intelligent Data Understanding (CIDU)*, pages 47–54, Oct 2012.
- [15] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 12 2014.
- [16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML'15, page 448–456. JMLR.org, 2015.
- [17] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [18] Lynton Ardizzone, Jakob Kruse, Sebastian J. Wirkert, Daniel Rahner, Eric W. Pellegrini, Ralf S. Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. *CoRR*, abs/1808.04730, 2018.
- [19] Raphael Shirley, Yannick Roehly, Peter D. Hurley, Veronique Buat, María del Carmen Campos Varillas, Steven Duivenvoorden, Kenneth J. Duncan, Andreas Efstathiou, Duncan Farrah, Eduardo González Solares, Katarzyna Malek, Lucia Marchetti, Ian McCheyne, Andreas Papadopoulos, Estelle Pons, Roberto Scipioni, Mattia Vaccari, and Seb Oliver. HELP: a catalogue of 170 million objects, selected at 0.36-4.5 μm , from 1270 deg^2 of prime extragalactic fields. *Monthly Notices of the Royal Astronomical Society*, 490(1):634–656, November 2019.