Astronomical Image Coaddition with Bundle-Adjusting Radiance Fields

Harlan Hutton
Flatiron Institute
New York University
ahh303@nyu.edu

Harshitha Palegar
Flatiron Institute
New York University
hpalegar98@gmail.com

Jenna Eubank
Flatiron Institute
New York University
jke261@nyu.edu

Shirley Ho
Flatiron Institute
Princeton University
shirleyho@flatironinstitute.org

Miles Cranmer
Princeton University
miles.cranmer@gmail.com

Peter Melchior
Princeton University
melchior@astro.princeton.edu

Abstract

Image coaddition is of critical importance to observational astronomy. This family of methods consisting of several processing steps such as image registration, resampling, deconvolution, and artifact removal is used to combine images into a single higher-quality image. An alternative to these methods that are built upon vectorized operations is the representation of an image function as a neural network, which has had considerable success in machine learning image processing applications. We propose a deep learning method employing gradient-based planar alignment with Bundle-Adjusting Radiance Fields (BARF) to combine, de-noise, and remove obstructions from observations of cosmological objects at different resolutions, seeing, and noise levels – tasks not currently possible within a single process in astronomy. We test our algorithm on artificial images of star clusters, demonstrating powerful artifact removal and de-noising.

1 Introduction

Combining images of celestial objects is a common and important task performed by astronomers. The goal of *image coaddition* is to resample the images, which corrects for any warping that occurred during the observation, and add them to produce a single deeper, more meaningful representation (AlSayyad [1]). The first task is therefore *image registration*, which is usually carried out by mapping the observed locations of stars to their locations in a common reference, followed by *resampling* to map all given images into that reference coordinate system. Thereafter, coaddition involves direct linear combination, point spread function (PSF) matching, or non-linear weighting.

Image coaddition can serve additional purposes besides the effective reduction of noise, in particular aiding *deconvolution*, i.e. the deblurring operation that reduces the effects of telescope optics and Earth's atmosphere. However, deconvolution of noisy images is an ill-defined operation (Starck et al. [10]). As a result, astronomers need to adopt a coaddition scheme in which there exists a trade-off between improvements in image sharpness and effective noise level. Advanced coaddition methods such as likelihood coaddition, which calculates the joint likelihood of models of the true sky, or

Machine Learning and the Physical Sciences workshop, NeurIPS 2022.

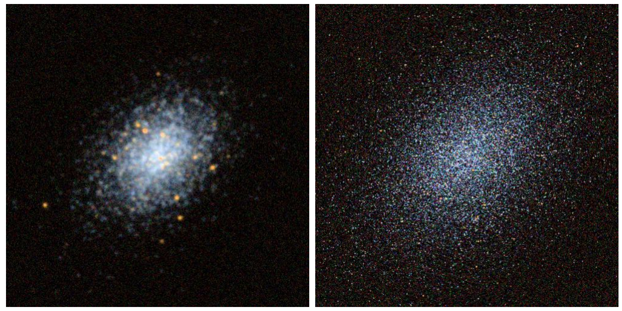

Figure 1: `ArtPop` Generated Low Resolution and High Resolution Images

decorrelated coaddition attempt to find that optimal trade-off but are computationally expensive and require specific conditions to operate (Bosch [2]). Another method, Kaiser coaddition, is less expensive but has strict demands for use: 1) noise in the images is uncorrelated and white, 2) the PSFs are spatially constant and 3) no pixels or boundaries are missing in the images (Kaiser [5]). In summary, coaddition is a longstanding objective in astronomy for which formally optimal solutions exist but with the caveat of needing idealized conditions.

The situation becomes even more complex when the goal becomes to combine images from different instruments. This task is becoming of increasing interest as several large astronomical surveys are being planned in order to observe large overlapping regions of cosmological scenes. Such an undertaking would benefit from large-scale image coaddition (Rhodes et al. [9]). Well-documented solutions exist to successfully combine images of different resolutions, but their common trade-off is accuracy of the resampled image versus computational efficiency.

In order to address the shortcomings of current image coaddition methods, we present a single, more streamlined image processing technique for all of the tasks listed above: Bundle-Adjusting Radiance Fields (BARF). BARF (Lin et al. [6]) builds off of the 3D computer vision scene reconstruction method Neural Radiance Fields (NeRF) (Mildenhall et al. [8]). BARF's key difference to common techniques in astronomy is that it represents the image not as a regular collection of pixels but as a neural network that maps sky positions to light intensity. The network can be trained robustly through the combination of images even if their respective warps are initially unknown. In this work we focus on the removal of artifacts when coadding images with different resolutions, seeing, and noise levels in order to show how planar alignment can perform coaddition, deconvolution, and artifact masking. Our code is available at `https://github.com/harlanhutton/AstroNeRF`.

2 Data

The python package `ArtPop` allows us to create artificially-rendered star clusters by manipulating object characteristics, instrument properties and capture quality (Greco and Danieli [3]). Outputs are either black and white, replicating a single bandpass filter, or colorized, combining multiple telescope filters. The process of generating an artificial image is broken into 3 steps: defining the source object, creating the imager, and observing. The source represents the star cluster, the imager represents the hypothetical telescope, and the observation is the act of using the imager to capture the source.

By manipulating the `ArtPop` parameters, we can mimic a single star cluster that has been captured from different telescopes under various environmental conditions. Perturbing the noise parameter allows us to produce a dataset with various levels of sky noise. To vary seeing levels, we choose different full-width half-maximum (FWHM) values for the PSF in the observation step. To vary resolutions, we choose different values of pixel scale.

Additionally, we simulate the complications in practical astronomical imaging, namely "warps", i.e. the non-linear mapping between image coordinates and sky coordinates due to e.g. the flexing of telescope, and partial obstructions from artifacts such as ghost or satellite trails. Warps are applied to crops of the original image, and obstructions are added to the warped image.

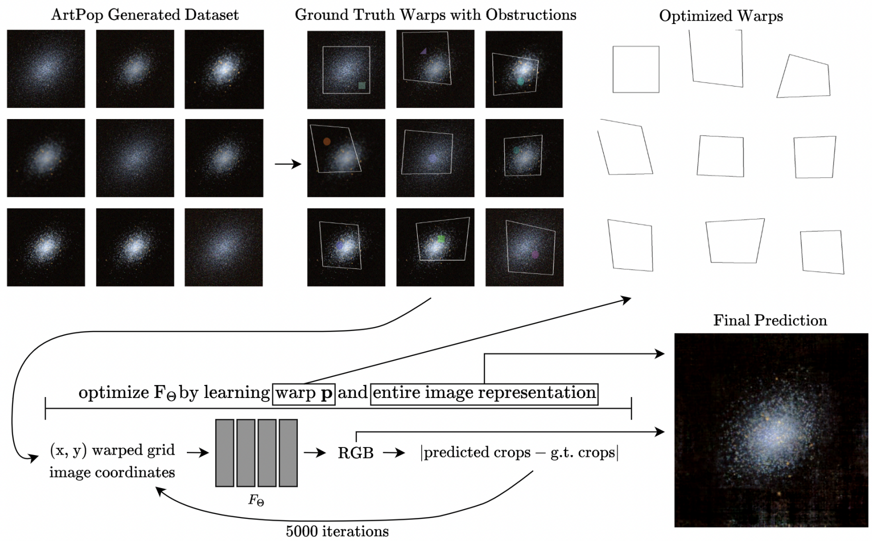

Figure 2: Diagram to show the workflow of planar alignment with BARF. The `ArtPop` images are generated. The images are cropped, ground truth warps are generated,and obstructions are added. The xy-coordinates of the ground truth warps act as input to the neural network which learns the warp parameters for each crop and the entire image representation used to make the final prediction.

3 Approach

The crucial change we make to BARF in order to adapt the algorithm to astronomy is to learn a single representation from a set of different input images rather than just one. Building off of Lin et al. [6], we seek to learn both an image representation and the warp parameters for each image. The key idea is to represent the image function $f : \mathbb{R}^2 \to \mathbb{R}^{\mathrm{B}}$, which maps from the 2D position $\mathbf{x}$ to the source intensity in B bands (e.g. $B = 3$ for RGB values), by a neural network. The objective function is

$$\min_{\mathbf{p}_{(1,\ldots,M)},\Theta} \sum_{i=1}^{M} \sum_{\mathbf{x}} |f(\mathcal{W}(\mathbf{x};\mathbf{p});\Theta) - \mathcal{I}_i(\mathbf{x})|, \tag{1}$$

where $\mathbf{p} \in \mathbb{R}^8$ are the warp parameters, $\mathcal{W}$ is the warp transformation, $\mathcal{I}$ is the image, Θ is the network parameters, and M is the number of images. We use the L1 loss to better account for the strong outliers random obstructions can cause. Through the optimization of Equation 1, we find the homography warp parameter for each patch while also learning the neural representation of the entire predicted image at arbitrary resolution.

Whereas Lin et al. [6] cropped a single ground truth image into multiple patches and applied warp to those patches, we crop different `ArtPop` images of the same star cluster. Each patch is then warped using translation and homography. Homography is parameterized with sl(3) Lie algebra which allows for a fixed scale ensuring only 8 free parameters are estimated. We justify using sl(3) as it accounts for when the observed plane does not strictly pass through the optical center of the camera (Mei et al. [7]). To each warped crop we then add random obstructions to demonstrate simultaneous coaddition and artifact removal.

4 Experiment

Using the same star cluster, we create a dataset of nine 500×500 pixel images randomly generated at high resolution or low resolution. Once the training images are created, we add randomly-generated obstructions that vary in shape, size, color, and location. The goal of adding these exaggerated obstructions is to demonstrate how planar alignment with BARF removes them from the final prediction. For each input image we generate a 300×300 cropped pixel patch with homography perturbation. All patches are initialized at the center and the first patch's warp is the identity so that

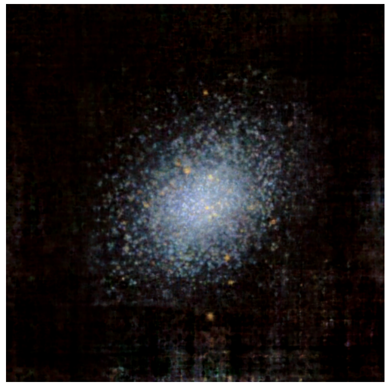

Figure 3: Final Prediction

the final prediction can implicitly align with the ground truth. We visualize the added obstructions and ground truth warps for each image in Figure 2.

To minimize the objective function and find the warp parameter for each patch and the neural representation of the entire image, we use the same network as Lin et al. [6]: a simple Multi-Layer Perceptron with a ReLU activation function and four 256-dimensional hidden units trained for 5000 iterations. We use the Adam optimizer with a learning rate of 0.001 for both the neural image representation and the warp parameters and a weight decay of 1×10^{-8} (Joseph et al. [4]).

5 Results

We visualize the final prediction image in Figure 3. We find that planar alignment with BARF successfully removes obstructions and reduces noise while learning the neural representation of the scene, even with input images at various resolutions, seeing, and noise levels. The prediction quality remains high even for a small batch size of 9.

In this context of coaddition with input images greatly varying in resolution, we consider the final prediction deconvolved overall and not just when compared to low quality images. We argue it is truly deconvolved because it corrects the systematic error of blur in the smaller features in the collective images. Although the final output might not be the optimal highest quality image of the stellar object, it still manages to preserve important information from each input image while reducing partial blur and removing obstructions when compiled.

As this experiment serves primarily as a demonstration of the capabilities of BARF, we do not list specific metrics of success, as they would inherently depend on the scientific purpose of the image coaddition. Instead, we demonstrate and focus on the visual quality of the coadded images.

6 Conclusion

We present planar alignment with BARF as a novel method to coadd, de-convolve, and remove obstructions from astronomical images, even when taken at different resolutions.

This deep learning strategy is simple and flexible when compared to current image processing methods in astronomy. Applying a neural network to implicit image position representations, instead of vector representations, generates an image function that has promising applications for further streamlining of current image processing techniques. Because coaddition, deconvolution, and artifact masking are longstanding objectives in astronomy with considerable research done to optimize each, we do not introduce planar alignment as a strategy to outperform any singular step in image processing. Instead, we show it as an alternate route to solving the multi-resolution problem that provides benefits of flexibility and simplicity.

One current shortcoming of this method is how prediction quality breaks down towards the edges of the image, as there is less information due to the cropped patches. Prediction quality also decreases as input image FWHM and noise levels are raised. However, even with a random obstruction in every input image, the output image shows the reconstructed star cluster with little trace of obstructions. Additionally, there is no current metric other than visual observation to compare performance to traditional methods. It is important to note that the scope of this paper only focuses on improvement of visual quality. Although outside the scope of this paper, we welcome suggestions for performance metrics of BARF that are meaningful in astronomy.

Future work aims to quantitatively compare the performance of planar alignment with BARF to traditional astronomical methods of image processing in terms of cost-effectiveness, luminosity, and overall scene reconstruction. Although an early experimentation, planar alignment with BARF shows that computer vision methods like homography estimation could prove to be the successful route for image processing. Increased technological advancements in telescope imaging may warrant the need for large-scale automation robust to noisy data that a traditional machine learning homography estimation method like planar alignment can provide.

Broader Impact Statement

In this paper we introduce BARF as a potential strategy to address the shortcomings of traditional image coaddition methods in astronomy. A neural approach could become a better, more efficient method for image coaddition in terms of simplicity and flexibility. This would allow astronomers to produce high-quality coadded images faster, with less data, and with lower quality data than traditional methods allow. Not only would this benefit the astronomy community, but also society as a whole, as it would aid in the technological advancement of space exploration. Additionally, this method is not limited to the realm of astronomical coaddition. A future avenue to explore would be the application of homography estimation methods like BARF to the coaddition of satellite images of Earth. The flexibility of the imaging function could be more robust than current methods to the removal of smoke or clouds when reconstructing scenes of landscapes.

References

[1] Yusra AlSayyad. Coaddition artifact rejection and comparewarp. 2019.

[2] Jim Bosch. Flavors of coadds. 2016.

[3] Johnny P. Greco and Shany Danieli. Artpop: A stellar population and image simulation python package. 2021.

[4] Rémy Joseph, Peter Melchior, and Fred Moolekamp. Joint survey processing: combined resampling and convolution for galaxy modelling and deblending, 2021. URL `https://arxiv.org/abs/2107.06984`.

[5] Nick Kaiser. Addition of images with varying seeing. 2001.

[6] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. *CoRR*, abs/2104.06405, 2021. URL `https://arxiv.org/abs/2104.06405`.

[7] Christopher Mei, Selim Benhimane, Ezio Malis, and Patrick Rives. Homography-based tracking for central catadioptric cameras. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 669–674. IEEE, 2006.

[8] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *CoRR*, abs/2003.08934, 2020. URL `https://arxiv.org/abs/2003.08934`.

[9] Jason Rhodes, Robert C Nichol, Éric Aubourg, Rachel Bean, Dominique Boutigny, Malcolm N Bremer, Peter Capak, Vincenzo Cardone, Benoît Carry, Christopher J Conselice, Andrew J Connolly, Jean-Charles Cuillandre, N A Hatch, George Helou, Shoubaneh Hemmati, Hendrik Hildebrandt, Renée Hložek, Lynne Jones, Steven Kahn, Alina Kiessling, Thomas Kitching, Robert Lupton, Rachel Mandelbaum, Katarina Markovic, Phil Marshall, Richard Massey, Ben J Maughan, Peter Melchior, Yannick Mellier, Jeffrey A Newman, Brant Robertson, Marc Sauvage, Tim Schrabback, Graham P Smith, Michael A Strauss, Andy Taylor, and Anja

Von Der Linden. Scientific synergy between LSST and euclid. *The Astrophysical Journal Supplement Series*, 233(2):21, dec 2017. ISSN 0067-0049. doi: 10.3847/1538-4365/aa96b0. URL `http://iopscience.iop.org/article/10.3847/1538-4365/aa96b0/meta`.

[10] J.L. Starck, E. Pantin, and F. Murtagh. Deconvolution in astronomy: A review. *Publications of the Astronomical Society of the Pacific*, 114(800):1051–1069, 2002. ISSN 00046280, 15383873. URL `http://www.jstor.org/stable/10.1086/342606`.

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)? [N/A]
   (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)? [N/A]
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? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (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]