Single Image Super-Resolution with Uncertainty Estimation for Lunar Satellite Images

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recently, there has been a renewed interest in returning to the Moon, with many planned missions targeting the south pole. This region is of high scientific and commercial interest, mostly due to the presence of water-ice and other volatiles which could enable our sustainable presence on the Moon and beyond. In order to plan safe and effective crewed and robotic missions, access to high-resolution (<0.5 m) surface imagery is critical. However, the overwhelming majority (99.7%) of existing images over the south pole have spatial resolutions >1 m. In order to obtain better images, the only currently available way is to launch a new satellite mission to the Moon with better equipment to gather more precise data. In this work we develop an alternative that can be used directly on previously gathered data and therefore saving a lot of resources. It consist of a single image super-resolution (SR) approach based on generative adversarial networks that is able to super-resolve existing images from 1 m to 0.5 m resolution, unlocking a large catalogue of images ($\sim$50,000) for a more accurate mission planning in the region of interest for the upcoming missions. We show that our enhanced images reveal previously unseen hazards such as small craters and boulders, allowing safer traverse planning. Our approach also includes uncertainty estimation, which allows mission planners to understand the reliability of the super-resolved images.

1 Introduction

The Moon has been receiving increasing interest in recent years. Many upcoming missions have been planned, for example in the frame of NASA's Artemis program which aims to put humans back on the Moon within this decade (Smith et al., 2020). A major goal is to achieve a permanent presence on the Moon, and for many of these missions the south pole is the main target. This region is of high scientific and commercial interest because it is expected to host water-ice and other resources which could enable our sustainable presence.

High-resolution (HR) satellite imagery is essential for planning such missions. For example, when planning human and rover traverses, mission planners need imagery with a resolution of $\sim$0.5m/px or better to detect relevant features at the rover-/astronaut-scale, such as craters and boulders (Robinson et al., 2010). This allows small hazards to be identified in advance, ensuring safe and efficient traverses. Over the past 12 years NASA's Lunar Reconnaissance Orbiter (LRO) mission has been capturing state-of-the-art images of the lunar surface using its Narrow Angle Camera (NAC) (Robinson et al., 2010), and their spatial resolutions typically vary in the range 0.5-2 m. However, at the south pole the overwhelming majority of these images have resolutions >1m. This is due to the inherent low lighting conditions. Over the south pole, most images (99.7%) are captured using what is known as "summed mode" operation of the NAC, which sums two adjacent pixels to improve SNR at the

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

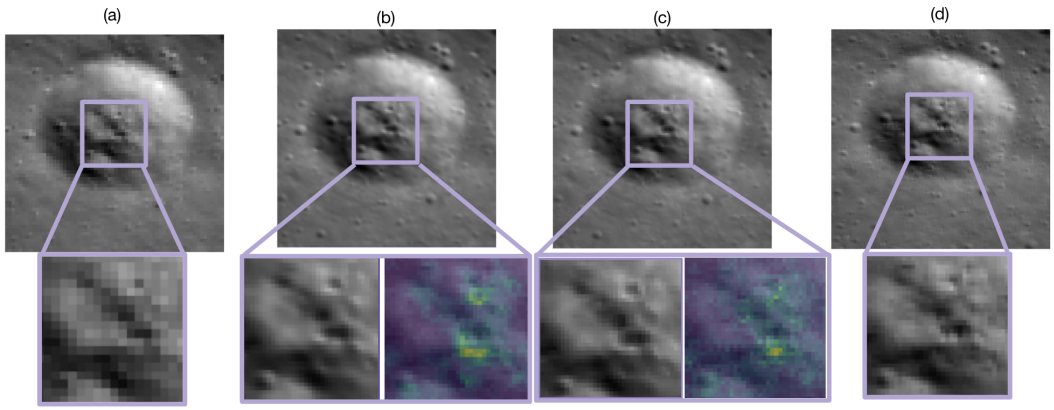

Figure 1: Using ANUBIS to super-resolve real LRO NAC imagery (Lat: -64.51, Long: 299.6). (a) A real 1 m/px summed mode image, (b) after LSRRes (our work), (c) after LSRGAN (our work) , (d) a coinciding real 0.5 m/px regular mode image (ground truth). The uncertainty map displayed next to the enhanced image shows our confidence in the super-resolved image (yellow: high uncertainty, purple: low uncertainty).

expense of reducing resolution by a factor of 2 (Humm et al., 2015). Unfortunately, this limits their usefulness for mission planning.

In this work we introduce ANUBIS (Adversarial Network for Uncertainty-Based Image Super-resolution). ANUBIS uses a deep ensemble of GANs to super-resolve these images, increasing their resolution by a factor of 2 (from 1 m to 0.5 m), providing enhanced data products for mission planners. We show that these images enable safer traverse planning, and by using an ensemble we are also able to output uncertainty estimates, which is essential for their safety-critical application.

2 Methodology

The main goal of this work is to learn to super-resolve existing summed-mode NAC images over the south pole from $\sim$1m/px to $\sim$ 0.5m/px. We note that a small number of "regular mode" images over the south pole are taken without the summing operation which have $\sim$0.5m/px resolution. These images are not common ($\sim$ 0.0028%), but there are enough examples with high SNR to construct a labelled training dataset. An example regular mode image is shown in Figure 1 (d).

Specifically, we generate a training, validation, and testing data starting from 121 (52,224$\times$5,064 px) regular mode NAC images of the south pole. The selected images have a resolution of 0.5 ± 0.05 m/px and similar illumination conditions to the summed mode polar images. We divide the images into 128×128 patches, resulting in 220,000 HR patches. To approximate the NAC summed mode, we apply a local 2×2 mean and round to the next integer value to obtain their low-resolution (LR) $\sim$1m/px counterparts. This operation represents a best effort approach to approximate the real summing performed on the spacecraft. We split this dataset into a training, validation and test dataset using a 80:10:10 split.

Two separate single image SR approaches are designed. The first uses a residual network (LSRRes) which takes a LR image and outputs a HR image, trained with an L2 loss function. The second builds on the first, training LSRRes with an additional adversarial loss (LSRGAN), and is shown in Figure 2. The goal of this approach is to ensure the output image is as realistic as possible via the interaction of the generator (LSRRes) and discriminator network.

A major concern when using single image SR is the potential for generating artifacts (or "hallucinations"), which are generally related to the ill-posedness of the inverse problem. This is of key concern in our application, as mission planners will use our images in safety-critical applications. Thus to increase the reliability of our images we also explore uncertainty estimation. Specifically, we train an ensemble of 24 LSRGANs with different weight initialisations, allowing us to output multiple realisations of the same super-resolved image. Each realisation differs slightly due to ill-posedness of the problem. The pixel-wise standard deviation of these images is used to generate an uncertainty

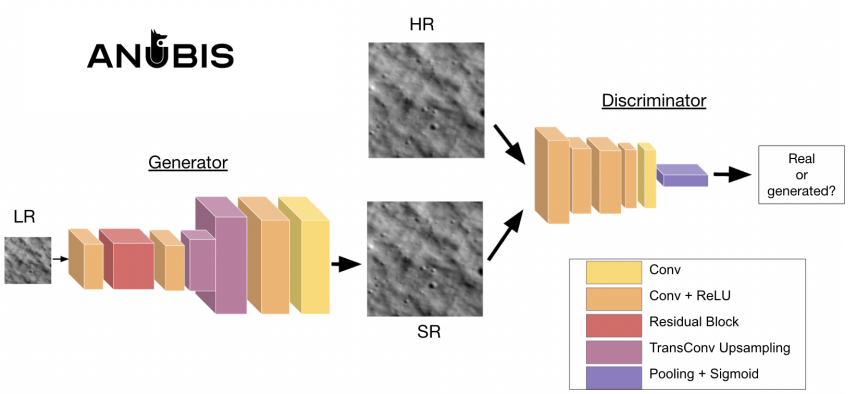

Figure 2: The ANUBIS LSRGAN architecture, consisting of a residual generator and a fully-convolutional discriminator. Colors indicate different block types.

Method	MAE↓	PSNR↑	SSIM↑
LR	0.0042	44.81	0.9745
Bilinear	0.0037	45.89	0.979
Bicubic	0.0032	47.26	0.984
LSRGAN (ours)	0.0034	47.23	0.983
LSRRes (ours)	**0.0027**	**49.13**	**0.989**

Table 1: Performance of SR methods on our test set. Best scores in bold. The LSRRes/LSRGAN metrics are computed using only a single randomly selected ensemble member.

map, which provides an estimate of the reliability of our SR process. The same uncertainty workflow is carried out for the LSRRes-only approach too. All training details are described in Appendix A.

3 Results and discussion

An example real 1 m/px summed mode image over the south pole is shown in Figure 1 (a), and the result after applying LSRRes and LSRGAN (using a single randomly selected ensemble member) is shown in (b) and (c). A real 0.5 m/px regular mode image is also available over this region, and is shown in (d). We find that both LSRRes and LSRGAN are able to significantly enhance the resolution of the LR image, revealing small craters and other surface features which are present in the HR image and difficult to see in the LR image. When evaluating more images, we find that LSRGAN tends to output images with higher perceptual quality than LSRRes. Specifically, LSRGAN adds more small scale features such as craters and boulders to the image compared to LSRRes, resulting in a frequency spectrum which is more comparable to ground truth HR images.

We evaluate the quantitative performance of our methods on our test set of images, shown in Table 1. As metrics, we consider the mean absolute error (MAE), peak signal to noise ratio (PSNR), and structural similarity index measurement (SSIM). We also compare our methods to two baseline approaches (bilinear and bicubic interpolation). Under these metrics, we find that LSRRes gives the highest performance. However, given the improved perceptual quality of the LSRGAN images observed above, these metrics may not be the best suited for this task.

We also show the uncertainty maps generated by our ensembles of LSRRes/LSRGANs for the super-resolved images in Figure 1. These maps are most uncertain around crater rims, which is believable as they represent high-frequency features in the image which are likely difficult to reconstruct.

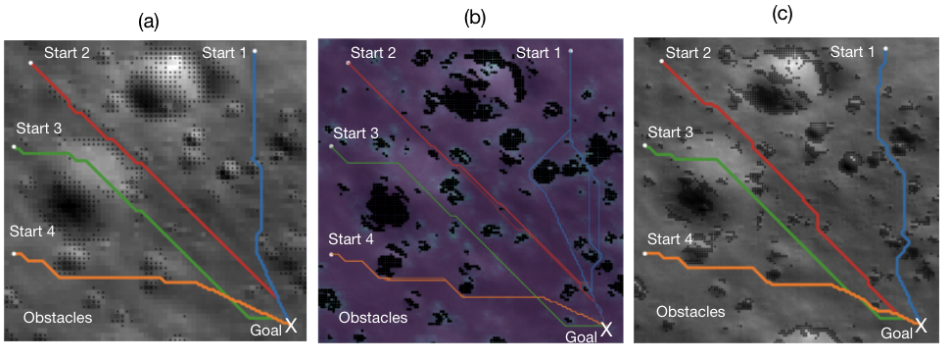

Figure 3: Validating ANUBIS with traverse planning. (a) Low-resolution (input), (b) Super-resolution + Uncertainty map (our work), (c) High-resolution (Ground truth).

4 Validation with traverse planning

Finally, we show the value of our super-resolved images on a downstream task, here: traverse planning. We set up a hypothetical problem, where a rover must plan a safe traverse from its starting position to a goal, given an image and the locations shown in Figure 3. Given an image, we detect obstacles based on the image gradients and then use an A* path planning algorithm to plan a safe traverse around them. During path planning using our super-resolved images, we also add obstacles derived from our uncertainty map (using the same obstacle detection process, i.e., high uncertainty pixels are considered as obstacles). Figure 3 shows the resulting traverses when using a LR summed mode image, its super-resolved version using LSRGAN, and a ground truth regular mode HR image. We find that for this location the LSRGAN paths more closely match the HR image paths than the LR image paths, i.e. allowing for a safer and more efficient traverse.

5 Limitations and future work

There are multiple limitations and future directions of our work. Firstly, we restricted ourselves to only learning a mapping from 1 m/px to 0.5 m/px, whereas in reality NAC images of the south pole span a range of resolutions >1 m/px. In future work, we aim to develop a workflow which is able to learn a mapping from any input resolution to any output resolution, within reason, allowing more NAC images to be super-resolved. Secondly, there are many other uncertainty estimation techniques which could be tested. For example, GANs with random input vectors (Abid et al., 2021), normalising flows (Lugmayr et al., 2020), dropout or other Bayesian inference approaches could be more effective. Finally, we assume that our downsampling operator can be approximated by a simple summation and rounding. However, in reality NAC images are subject to a lossy compression before downlink to Earth (Robinson et al., 2010), and therefore this degradation process is not modelled; we intend to include this in future work.

6 Conclusions and broader impact

We have investigated methods for super-resolving lunar imagery over the south pole by a factor of 2 (from 1 m/px to 0.5 m/px) and shown that they can improve image quality and help achieve downstream tasks such as traverse planning. Furthermore we provided uncertainty estimates for each super-resolved image, which is essential for safety-critical applications. Our approach provides higher resolution images over the south pole of the Moon, where they are not currently available, allowing mission planners to plan safer and more effective missions.

This work could have a positive impact on the planning and execution of future lunar exploration missions, specifically by reducing their risk and maximizing their efficacy. The uncertainty maps that are delivered along with the improved images greatly increase the reliability of ANUBIS for safety-critical applications. Upcoming robotic and crewed missions, such as NASA's VIPER or other Artemis missions, are a few of many possible beneficiaries of our work.

References

Abid, M. A., Hedhli, I., and Gagné, C. (2021). A generative model for hallucinating diverse versions of super resolution images. *arXiv preprint arXiv:2102.06624*.

Humm, D., Tschimmel, M., Brylow, S., Mahanti, P., Tran, T., Braden, S., Wiseman, S., Danton, J., Eliason, E., and Robinson, M. (2015). Flight calibration of the lroc narrow angle camera. *Space Science Reviews*, 200.

Lugmayr, A., Danelljan, M., Van Gool, L., and Timofte, R. (2020). Srflow: Learning the super-resolution space with normalizing flow. In *European Conference on Computer Vision*, pages 715–732. Springer.

Robinson, M., Brylow, S., Tschimmel, M., Humm, D., Lawrence, S., Thomas, P., Denevi, B., Bowman-Cisneros, E., Zerr, J., Ravine, M., et al. (2010). Lunar reconnaissance orbiter camera (lroc) instrument overview. *Space science reviews*, 150(1-4):81–124.

Smith, M., Craig, D., Herrmann, N., Mahoney, E., Krezel, J., McIntyre, N., and Goodliff, K. (2020). The artemis program: An overview of nasa's activities to return humans to the moon. In *2020 IEEE Aerospace Conference*, pages 1–10. IEEE.

A Training details

Our LSRRes network uses a convolutional residual architecture with 6 layers, and ReLU activation functions (Generator section in Fig. 2). The discriminator in our LSRGAN uses a convolutional architecture with 6 layers, and ReLU activation functions. All approaches were trained using the Adam optimizer, a learning rate of 0.0002 and a batch size of 256. The inputs and outputs of the networks are normalised before training. Training takes approximately 3 hours on a single NVIDIA Tesla A100 GPU.

To create the deep ensemble of GANs used to generate the SR image distribution, the same architecture was trained 24 different times using random weight initialization.

B NeurIPS 2021 Paper Checklist

1. *Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?:* Yes
2. *Have you read the ethics review guidelines and ensured that your paper conforms to them?:* Yes
3. *Did you discuss any potential negative societal impacts of your work?:* Yes
4. *Did you describe the limitations of your work?:* Yes
5. *Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)?:* The data used during this work is publicly available at `https://www.lroc.asu.edu/`. The code will be released in a future journal publication.
6. *Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)?:* Yes
7. *Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)?:* Yes, with the uncertainty estimation.
8. *Did you include the amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)?:* Yes