

Quality-Weighted Iterative Deconvolution (QWID)

Michael Werth, Trent Kyono, Jacob Lucas

The Boeing Company

Justin Fletcher

Odyssey Consulting

Ian McQuaid, Michael Brannon

Air Force Research Laboratory

ABSTRACT

Multi-frame blind deconvolution (MFBD) techniques are often developed without foreknowledge of the incoming data quality. If a multi-frame ensemble contains a mix of high- and low-quality images, the image reconstruction quality may be worse than if only the high-quality images are processed. We report on a new MFBD algorithm that uses image quality scores to reduce processing time and improve reconstruction quality. These scores are automatically provided by a neural network that has been trained to assess the quality of ground-based observations of LEO satellites. This implementation, titled Quality-Weighted Iterative Deconvolution (QWID), is demonstrated on realistic ground-based observations.

1. INTRODUCTION

Observing a Low Earth Orbit (LEO) satellite with a large ground-based aperture has the potential to produce high-resolution images, but turbulence significantly limits the effective aperture size of such systems [1]. Mitigation of turbulence and other sources of noise can be attempted with various post-processing algorithms that range in complexity, from simpler techniques broadly described as lucky imaging [2, 3] to computationally-expensive blind deconvolution algorithms such as Likelihood-based Uncertainty-Constrained Iterative Deconvolution (LUCID) [4, 5], Physically-Constrained Iterative Deconvolution (PCID) [6, 7, 8], and Bispectrum [9, 10], among many others. One common challenge for post-processing approaches like these is that the input image frames are expected to conform to a number of unstated assumptions. Some of these assumptions often include:

1. The target object is observed during pristine photometric sky conditions (e.g. no clouds);
2. The target object is well-centered on the focal plane array (FPA), and does not fall off of the edge of the FPA;
3. The light beam path does not become blocked for some systematic reason, such as due to a rotating filter wheel.

It's possible for any of these assumptions to be violated during actual data collection. For instance, the LEO could be observed while it transits into a region of the sky that is obscured by clouds, the track loop could fail and allow part of all of the object to drift off of the FPA, some other science need could force a sudden change in photometric filter, etc. These changes can introduce large numbers of frames where the SNR is unsuitably low for MFBD processing, corrupting some frame ensembles with frames that may only contain noise.

In this paper we describe a new MFBD implementation titled Quality-Weighted Iterative Deconvolution (QWID). It is a system which uses a Convolutional Neural Network (CNN) to assign quality scores to input images prior to running the LUCID implementation of MFBD, which has been modified to throw away low-scoring frames from automatically-generated ensembles. With the threshold set sufficiently low, this results in a significant reduction in processing time without sacrificing reconstruction quality.

In Section 2 we briefly summarize the LUCID algorithm and the techniques it uses to rapidly reconstruct images from ensembles of multiple input frames as well as the CNN that scores input images. Section 3 describes the QWID algorithm, which incorporates the CNN-based scores for preconditioning input frame ensembles. In Sections 4 and 5 some of the performance improvements of QWID related to LUCID are characterized. The conclusions drawn from this effort are discussed in Section 6.

**DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048**

2. BACKGROUND

QWID's development was contingent on two previously-published efforts: LUCID and Automated Quality Scoring.

2.1 LUCID

LUCID is an implementation of MFBD [11, 12, 13] that assumes a simple linear forward model for the process of imaging through the atmosphere from a ground-based observatory. In a collection of $m = 1 \dots M$ short-exposure measurement frames, each m th measurement of image data, $d_m(\mathbf{x})$, represents a pristine image of the object, $o(\mathbf{x})$, convolved with a Point Spread Function (PSF) approximating the complete optical disturbance, $h_m(\mathbf{x})$, plus an additive noise realization term that includes Gaussian and Poisson sources, $n_m(\mathbf{x})$. This imaging model is described in the Equation 1:

$$d_m(\mathbf{x}) = o(\mathbf{x}) * h_m(\mathbf{x}) + n_m(\mathbf{x}) \quad (1)$$

In practice, Equation 1 may only hold for a few seconds in less when a stationary observer watches a LEO object. Thus, longer collections are split into numerous short ensembles that are processed independently, resulting in an embarrassingly parallel computational problem. Processing each ensemble entails simultaneously estimating the spatially-invariant object ($o_{est}(\mathbf{x})$) and the PSFs ($h_{est_m}(\mathbf{x})$) that are varying between each measurement frame by utilizing a maximum likelihood estimation (MLE) approach [14]. LUCID attempts to iteratively minimize the cost function shown in Equation 2:

$$J[o_{est}(\mathbf{x}), h_{est_1}(\mathbf{x}), h_{est_2}(\mathbf{x}), \dots, h_{est_m}(\mathbf{x})] = \sum_{m=1}^M \sum_{n=1}^N \frac{(i_m(\mathbf{x}_n) - d_m(\mathbf{x}_n))^2}{\sigma^2(\mathbf{x}_n) + i_m(\mathbf{x}_n)} \quad (2)$$

Equation 2 is a summation over M frames containing N pixels each. The numerator of this sum contains the noise-normalized squared residual between the measured data frames $d(\mathbf{x}_n)$ and the data model $i(\mathbf{x}_n)$. The denominator contains Gaussian read noise variance per-pixel of $\sigma^2(\mathbf{x}_n)$ and Poisson shot noise variance of $i_m(\mathbf{x}_n)$. Each m th model of the data is calculated from the convolution of the estimated object frame with each estimated PSF:

$$i_m(\mathbf{x}_n) = o_{est}(\mathbf{x}_n) * h_{est_m}(\mathbf{x}_n) \quad (3)$$

LUCID is implemented in NVidia's Compute Unified Device Architecture (CUDA) [15, 16] and C++. It uses a Limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) quasi-Newton method [17, 18] implemented in CUDA to iteratively minimize Equation 2.

2.2 Automated Quality Scoring

Several related works have explored the applications of CNNs and deep learning to astronomy. For noise reduction, [19] recently presented a proof-of-concept neural network for denoising the bispectrum for astronomical image recovery on synthetic data. For classification, [20, 21, 22] investigated the application of object classification using neural networks on photometric light curves and showed promising results. Using Generative Adversarial Networks (GANs), [23] recovered features from artificially degraded images with worse seeing and higher noise than the original with a performance that far exceeded the capabilities of simple deconvolution. Additionally, [24] used a GAN to generate more realistic images of galaxies than existing state of the art. [25] used machine learning to automatically segment and label galaxies in astronomical images. [26] showed promising results using an autoencoder for real-time MFBD of solar images. Additionally there have been probes into image scoring with deep learning; [27] applies a CNN to images and yields a 'human opinion' quality score, and [28] uses a relatively shallow network to give a quality score to distorted images. Later [29] applied CNNs to predict Space-object National Imagery Interpretability Rating Scale (SNIIRS) [30, 31] for ground-based images of satellites under various atmospheric perturbations. Lastly, [32, 33] attempt to determine the likelihood of whether a perturbed image will result in an coherent MFBD reconstruction. The success of these approaches motivates our application of similar networks to the scoring of ground based LEO (Low Earth Orbit) observations.

We desire an autonomous scoring model according to the following formalization. Let X be a set of simulated astronomical images corresponding to a collection, and Y be the collections score. We refer to each collection $x_i \in X$ as

**DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048**

a pass containing n sequential images. Given a pass $x_i \in X$, our primary design goal is to train a network $f : X \rightarrow Y$, which takes as input a collection $x_i \in X$ and provides a scoring prediction, $y_i \in Y$. To do this we train a single image classifier $g : X_j \rightarrow Y_j$ that takes as input a single image, $x_{ij} \in X_j$, for the j^{th} image belonging to a pass x_i and makes a prediction $y_{ij} \in Y_j$ for image x_{ij} . We then repeat this action for every image in pass x_i .

We used the prior scoring method outlined in [34, 29]. We use their best performing network InceptionResNetV2 (IRV2) [35]. We replaced the top dense layers with the following sequence of top layers (non-convolutional layers), a global average pooling layer, a 1024 neuron dense layer (*ReLU* activation followed by a dropout layer with rate of 0.2), a 256 neuron dense layer (*ReLU* activation followed by a dropout layer with rate of 0.2), and finally a 1 neuron dense layer with *linear* activation. Each dense layer was initialized with a Glorot normal distribution [36].

We trained IRV2 on the Simulated Images of Low-earth Objects (SILO) dataset [37], which provides scored SNIIRS labels for several million perturbed ground-based astronomical images. For input image augmentation, we applied random augmentation to each training image with the following specification: rotation within ± 30 degrees, horizontal flips, vertical flips, and zoom within $\pm 20\%$. Image augmentation was not used during testing inference. We used an iterative deepening training regime, where we first train the top dense layers by freezing all convolutional layers for 20 epochs. We then freeze the top dense layers and “unfreeze” the base model and train for 30 epochs. This is followed by another round of 20 epochs with base model frozen and the top dense layers trainable. We use the *Adam* optimizer with a learning rate of 10^{-4} for the first two iterations and 10^{-5} for the final training round. We used an MAE loss function, and saved IRV2 at the point with which the validation loss did not improve.

We use CLAHE augmentation, which is a method shown in [38] to improve classification performance particularly when issuing predictions on grayscale images on pre-trained ImageNet models that are optimized for 8-bit red-green-blue (RGB) inputs. Figure 1 shows the distribution of quality score network predictions for IRV2 versus truth scores for all images from the test set, i.e. objects never before seen by the network.

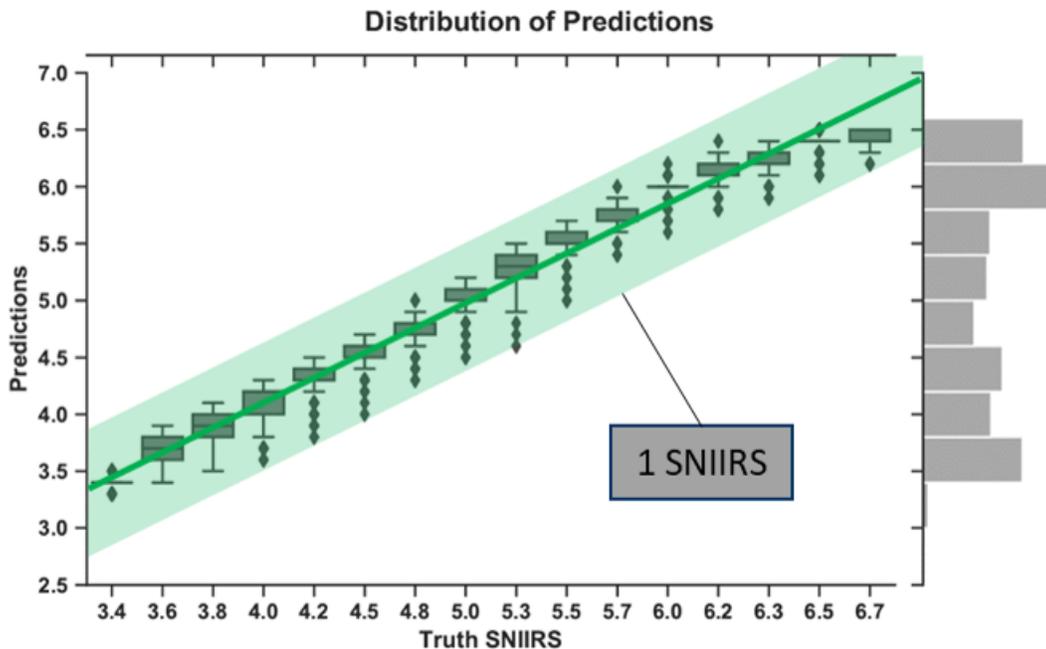


Fig. 1: Distribution of 2560 test set predictions from IRV2 evenly sampled over the full range of SNIIRS scores. Histogram at right displays density of predictions. A fitted mean and a shaded band with height of 1 SNIIRS around the mean, representing a standard uncertainty for a human analyst’s estimation of interpretability, are overlaid for comparison.

**DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048**

3. QUALITY-WEIGHTED ITERATIVE DECONVOLUTION: QWID

With the ability to automatically generate image quality scores described in Section 2.2, QWID is implemented as two modifications to the LUCID algorithm described in Section 2.1. These two modes are described as "ensemble-dropping" and "frame-dropping", respectively:

1. During ensemble generation, drop ensembles that contain $> 90\%$ frames with quality below some threshold;
2. During ensemble generation, drop frames from ensembles whose quality falls below some threshold.

These thresholds can be set to a very low value, one where QWID only rejects frames containing almost no signal. This is equivalent to using a simple signal to noise ratio (SNR) threshold. A slightly higher score threshold can take advantage of the scoring CNN's ability to accurately determine whether images are unsuitable for processing despite otherwise meeting an SNR threshold, such as when the object partially falls onto the edge of the FPA. Per the SNIIRS scale, an image with a score less than 1 is indistinguishable from noise, and less than 2 is indistinguishable from a point source. The CNN is trained to assign a score of 1 to frames where at least 10% of the object has fallen off of the FPA.

Dropping ensembles (item 1) should have the effect of improving overall processing time by virtue of simply not attempting to reconstruct an object when most of the ensemble's input frames fail to meet the MFBD assumptions outlined in Section 1. Dropping frames (item 2) should have the same effects as item 1 while also improving processing time and reconstruction quality in those ensembles where a small subset of the frames violate the MFBD assumptions.

4. PERFORMANCE EXPERIMENTS

Tests comparing the performance of LUCID and the two modes of QWID described in Section 3 were performed on commercial off-the-shelf (COTS) hardware. This system consists of one NVidia Tesla V100 [39] and one Intel Xeon E5-2698v4 CPU [40] coordinating GPU tasks. The GPU runs the scoring network described in Section 2.2 as a preprocessing step prior to launching QWID.

Three objects are selected from the SILO dataset: ARGOS (SCN 25634), HST (SCN 20580), and NOAA 8 (SCN 13923). Pristine renderings of these objects are shown in Figure 2.

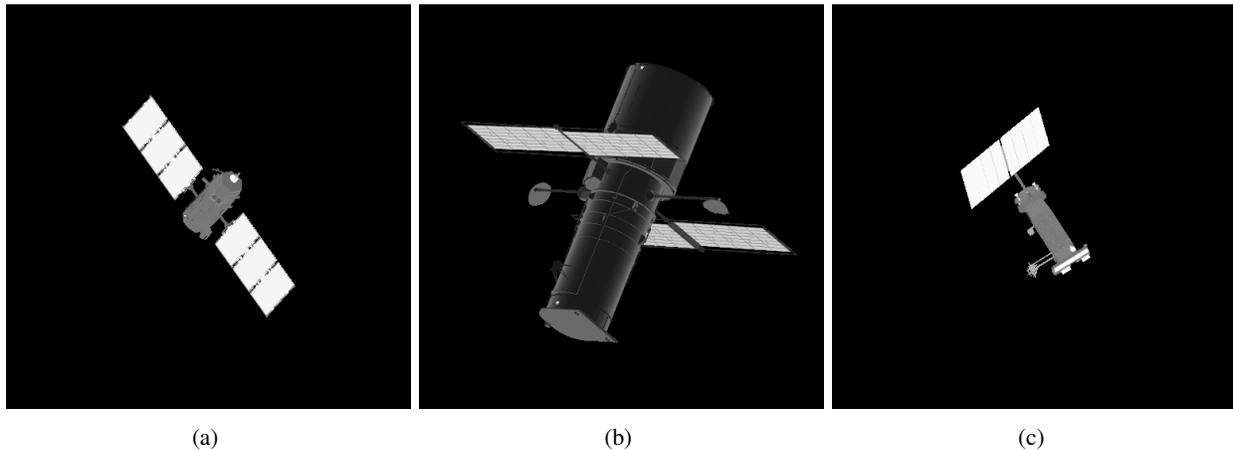


Fig. 2: Renderings of (a) ARGOS, (b) HST, and (c) NOAA 8 that were used for simulation.

The SILO dataset contains images of these objects observed under a standard set of observation scenarios, but the only images considered for this study consist of Komogorov observations of each target from a 3.6m telescope on Mt. Haleakala at $D/r_0 = 8$, at an altitude of 600 km and an observing range of 800 km. The renderings are downsampled from 512x512 to a 256x256 CCD camera with 8 ms exposures and 1.5 pe- of read noise. Each of these three scenarios

DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048

(e.g. three distinct objects) contain a total collection duration of 10,000 frames. Example images chosen from these scenarios are shown in Figure 3. In every 3,000 frame segment, the last 600 frames are replaced with Gaussian noise matching the camera’s read noise (1.5 pe-), representing frames where the object is no longer on the FPA for whatever reason; each scenario thus contains 1,800 noise-only frames and 8,200 frames containing the target.

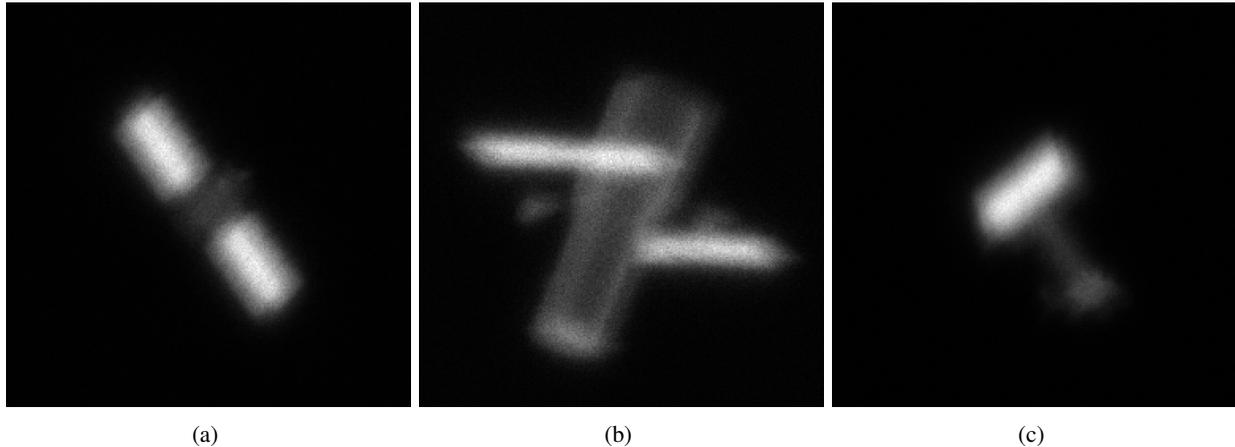


Fig. 3: Example images taken from datasets generated for QWID performance testing, consisting of observations of (a) ARGOS, (b) HST, and (c) NOAA 8 at $D/r_0 = 8$.

These datasets have their interpretability assessed by the scoring network, and then they are processed with LUCID, QWID using ensemble-dropping due to low quality, and QWID using frame-dropping due to low quality. The thresholds for quality are set to 1, corresponding to a SNIIRS score for an unresolved point source or noise-only frame. The time to score the interpretability of all images in each dataset is measured separately from the time to run QWID in its two modes.

5. RESULTS

Figure 4 shows the scoring network predictions for each of the 6 datasets. The scoring network accurately predicts that the frames containing no object have a very low score, whereas the score is much higher for frames that contain an atmospherically-degraded object. This confirms that a quality score threshold of 1 should exclude all noise-only frames from each dataset.

Noise-only frames cause significant issues in the reconstruction process. The highest-quality images from each dataset, created from ensembles that do not have any noise-only frames are shown in Figure 5. Since LUCID and QWID use the same underlying MFBD implementation, the highest-quality images are the same for each of the three techniques described in Section 4 (LUCID, QWID with ensemble-dropping, QWID with frame-dropping). The reconstructions shown in Figure 6 contain 120 input frames, but 60 are noise-only frames. These images have noticeable artifacts that degrade the resolution of each output image, despite half of the ensemble consisting of high-SNR frames. Ensembles that consist entirely of noise-only frames do not produce meaningful reconstructions and are not worth showing.

The first implementation of QWID described in Section 3 drops ensembles in which $> 90\%$ of the input frames fall below a SNIIRS threshold of 1. The second implementation goes a step further and drops low-scoring frames from ensembles, resulting in superior reconstructions. For instance, the same ensembles that produced low-quality reconstructions with LUCID (Figure 6) produce high-quality reconstructions with QWID’s second mode, as shown in Figure 7. These QWID reconstructions are almost indistinguishable from the reconstructions in Figure 5, in which the ensemble contained no noise-only frames.

Each of QWID’s two modes result in a reduction in MFBD processing time relative to LUCID processing time. However, the time that it takes to score all 10,000 frames can be significant. Figures 8- 10 show the ratio of (LUCID processing time) to (QWID processing time) for each scenario as a function of number of MFBD iterations.

**DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048**

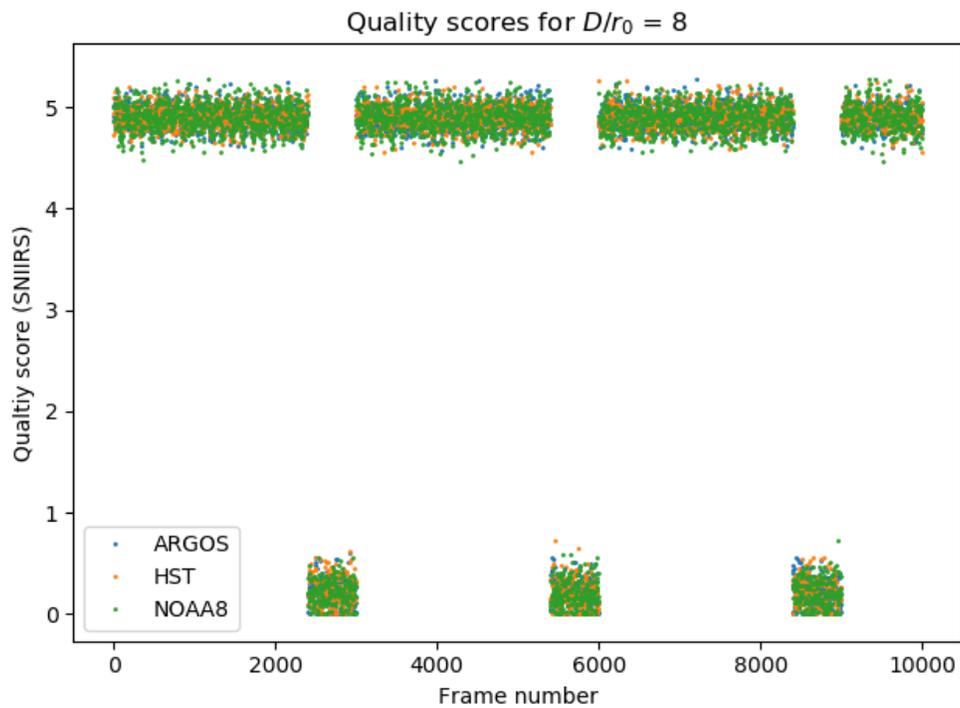


Fig. 4: Interpretability scores assigned by the CNN for ARGOS, HST, and NOAA 8 at $D/r_0 = 8$.

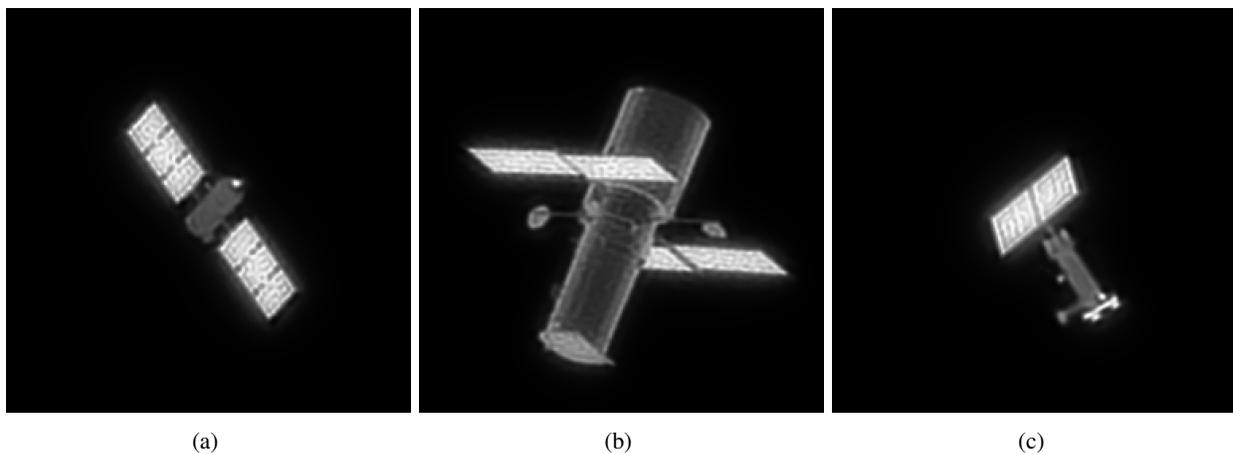


Fig. 5: Example LUCID reconstructions for (a) ARGOS, (b) HST, and (c) NOAA 8 at $D/r_0 = 8$ in ensembles that do not contain noise-only frames

DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048

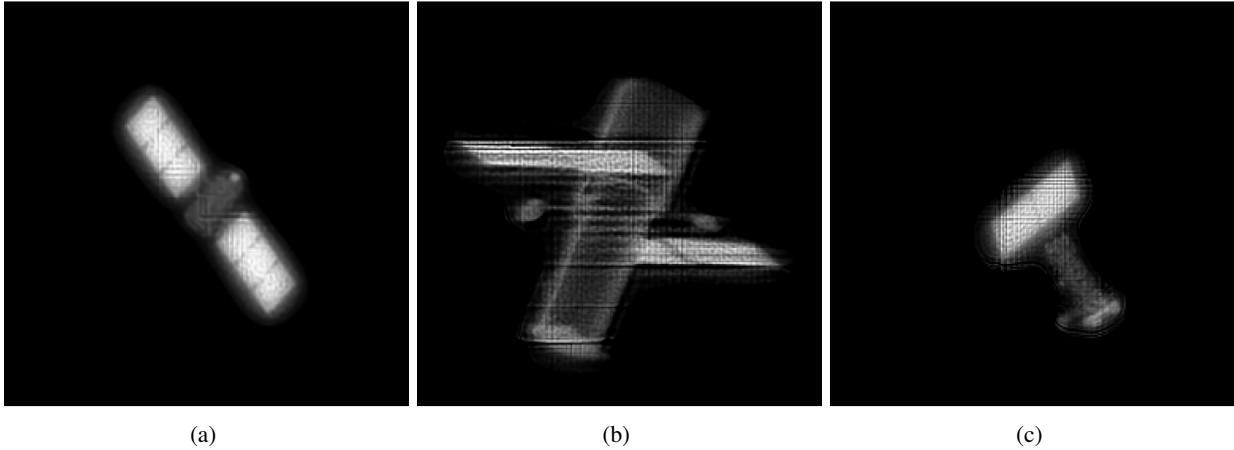


Fig. 6: Example LUCID reconstructions for (a) ARGOS, (b) HST, and (c) NOAA 8 at $D/r_0 = 8$ in an ensemble in which 50% of the frames only contain noise (60 frames of 120).

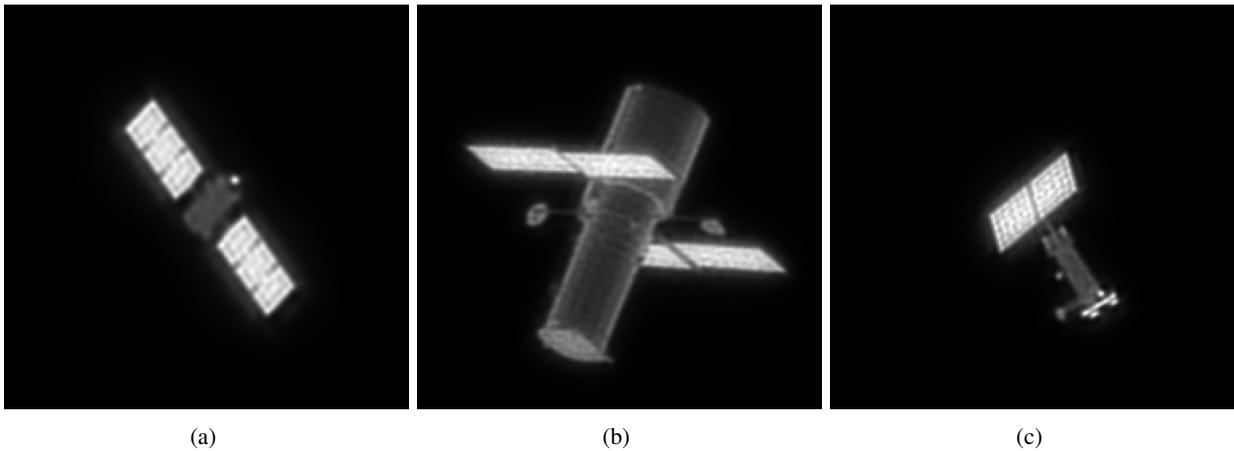


Fig. 7: Example QWID mode 2 reconstructions for (a) ARGOS, (b) HST, and (c) NOAA 8 at $D/r_0 = 8$ in an ensemble in which 50% of the frames only contain noise (60 frames of 120). These reconstructions were produced from the same input frames as the reconstructions in Figure 6

**DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048**

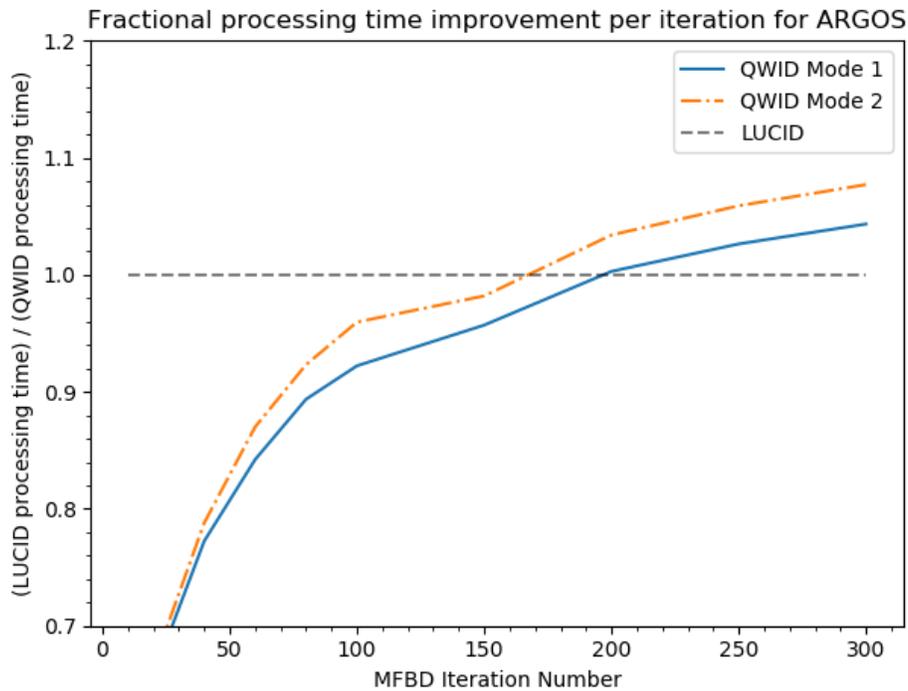


Fig. 8: Ratio of (LUCID processing time) to (QWID processing time) for each QWID’s two modes for ARGOS. A horizontal dashed line at 1.0 for reference; when a curve exceeds 1.0, total QWID processing time is that many times faster than LUCID.

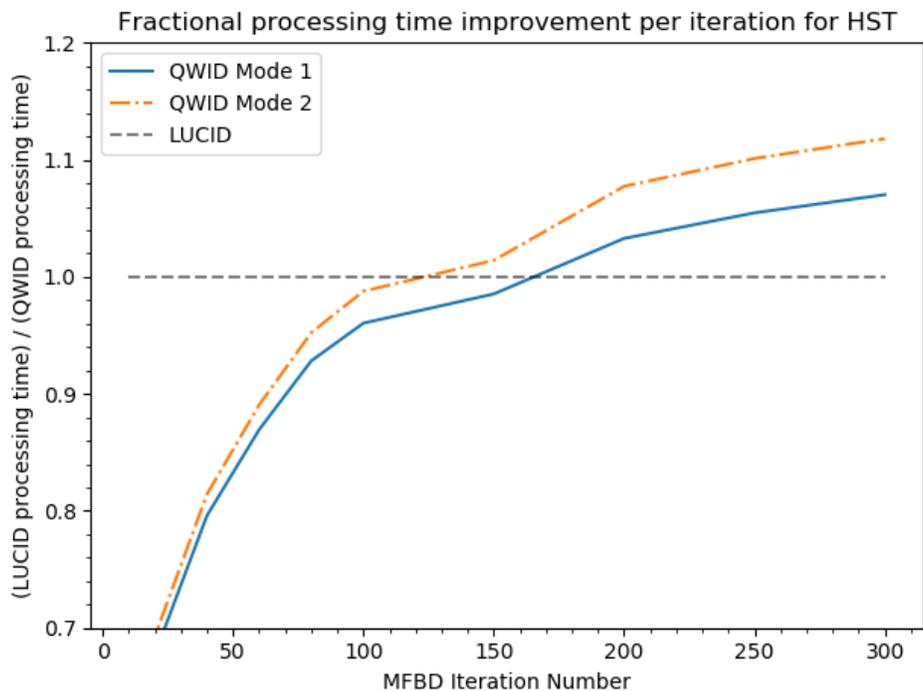


Fig. 9: Ratio of (LUCID processing time) to (QWID processing time) for each QWID’s two modes for HST. A horizontal dashed line at 1.0 for reference; when a curve exceeds 1.0, total QWID processing time is that many times faster than LUCID.

DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048

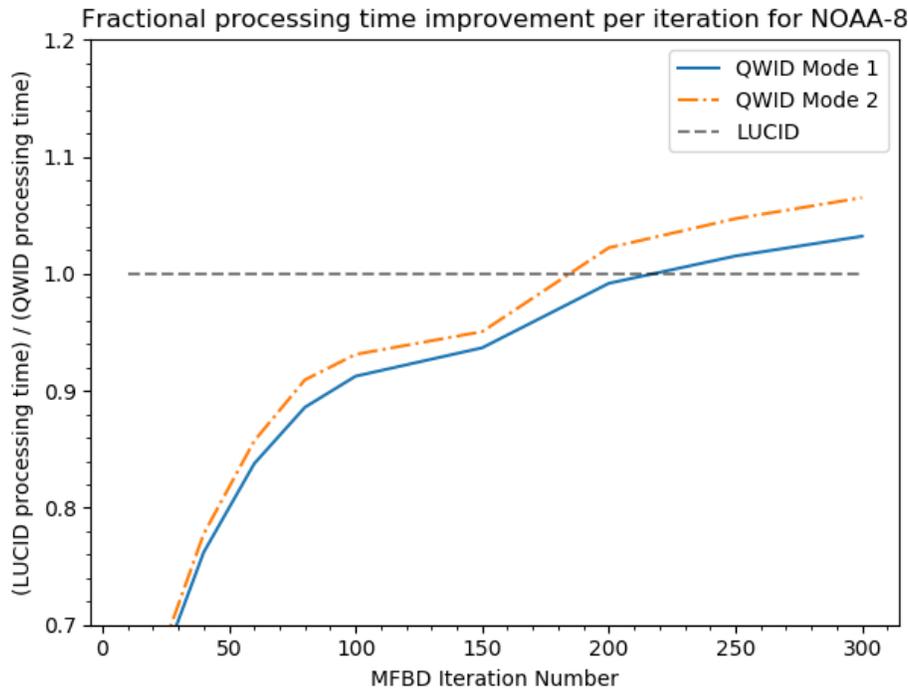


Fig. 10: Ratio of (LUCID processing time) to (QWID processing time) for each QWID’s two modes for NOAA 8. A horizontal dashed line at 1.0 for reference; when a curve exceeds 1.0, total QWID processing time is that many times faster than LUCID.

6. CONCLUSIONS

We have described QWID, an MFBD implementation that uses a CNN that automatically estimates the quality of input images in order to produce better MFBD reconstructions in less time. This is achieved by dropping low-scoring frames from the ensemble generation process, accomplishing removal of frame data that violates some of the numerous assumptions underlying many MFBD implementations, such as well-centered observations or photometric observing conditions.

Simulated passes containing a mix of 1800 noise-only frames and 8200 images with $D/r_0 = 8$ (e.g. a large aperture telescope with an AO system) were used for a series of performance tests. It was shown that MFBD ensembles consisting of 50% noise-only frames resulted in reconstructions with much worse quality than ensembles containing only target frames, as expected, and that using QWID’s frame-dropping mode resulted in reconstructions that were indistinguishable from ensembles containing no poor-quality frames. Both of QWID’s modes also conveyed improvements in processing time, but only for large numbers of iterations (> 150).

There are a number of additional improvements that can be made to QWID. First, the rationale for an ideal quality threshold has yet to be developed; the work described in this paper uses a very simple threshold based on limited empirical observations from the testing data. Second, the image quality scores could be incorporated into the cost function described in Equation 2. This may achieve improvements in image quality in ensembles where some frames have low average quality that does not fall below the threshold for frame removal. This study also only examined high-SNR AO-compensated images.

7. ACKNOWLEDGEMENTS

The authors would like to thank the Air Force Research Laboratory (AFRL) for sponsoring this research. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the

DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048

official policies, either expressed or implied, of the United States Air Force or the U.S. Government.

REFERENCES

- [1] Michael C. Roggeman and Byron M. Welsh. *Imaging Through Turbulence, Second Edition*. CRC Press, Mar 1996.
- [2] A. Smith, J. Bailey, J. H. Hough, and S. Lee. An investigation of lucky imaging techniques. *Monthly Notices of the Royal Astronomical Society*, 398(4):2069–2073, Oct 2009.
- [3] Wolfgang Brandner and Felix Hormuth. *Lucky Imaging in Astronomy*. Springer International Publishing, 2016.
- [4] Michael Werth, Brandoch Calef, and Kevin Roe. Multi-frame blind deconvolution accelerated with graphical processing units. In *AMOS Conference*, Sep 2019.
- [5] Michael Werth, Brandoch Calef, Kevin Roe, and Amanda Conti. Lucid: Accelerating image reconstructions of leo satellites using gpus. In *IEEE Aerospace Conference*, Mar 2020.
- [6] Julian C. Christou, E. Keith Hege, Stuart M. Jefferies, and Matthew Cheselka. Physically constrained iterative deconvolution of adaptive optics images. In *Proceedings Volume 3494, Atmospheric Propagation, Adaptive Systems, and Lidar Techniques for Remote Sensing II*, volume 3494, 1998.
- [7] Charles L. Matson, Kathy Borelli, Stuart Jefferies, Charles C. Beckner Jr., E. Keith Hege, and Michael Lloyd-Hart. Fast and optimal multiframe blind deconvolution algorithm for high-resolution ground-based imaging of space objects. *Appl. Opt.*, 48(1):A75–A92, Jan 2009.
- [8] Michael Werth, Brandoch Calef, Daniel Thompson, Kathy Borelli, and Lisa Thompson. Recent improvements in advanced automated post-processing at the amos observatories. In *IEEE Aerospace Conference*, pages 1–7, March 2015.
- [9] Pedro Negrete-Regagnon. Practical aspects of image recovery by means of the bispectrum. *J. Opt. Soc. Am. A*, 13(7):1557–1576, Jul 1996.
- [10] Charles L. Matson, Tom Soo Hoo, Maria Murphy, Brandoch Calef, Charles C. Beckner, Shiho You, Ron Vilorio, and Kathy Borelli. PCID and ASPIRE 2.0 - The Next Generation of AMOS Image Processing Software. *AMOS Conference Technical Proceedings*, page E60, 2007.
- [11] Timothy J. Schulz. Multiframe blind deconvolution of astronomical images. *J. Opt. Soc. Am. A*, 10(5):1064–1073, May 1993.
- [12] David Gerwe, Jim Stone, Carlos Luna, and Brandoch Calef. Comparison of maximum-likelihood image and wavefront reconstruction using conventional image, phase diversity, and lenslet diversity data. In *Unconventional Imaging II Proceedings*, volume 6307, Sep 2006.
- [13] Hirsch, M., Harmeling, S., Sra, S., and Schölkopf, B. Online multi-frame blind deconvolution with super-resolution and saturation correction. *A&A*, 531:A9, Feb 2011.
- [14] Timothy J. Holmes. Blind deconvolution of quantum-limited incoherent imagery: maximum-likelihood approach. *J. Opt. Soc. Am. A*, 9(7):1052–1061, Jul 1992.
- [15] John Nickolls, Ian Buck, Michael Garland, and Kevin Skadron. Scalable parallel programming with cuda. *Queue*, 6(2):40–53, March 2008.
- [16] NVIDIA Corporation. Cuda c programming guide v10.1.168, May 2019.
- [17] Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.*, 16(5):1190–1208, 1995.
- [18] Ciyou Zhu, Richard H. Byrd, Peihuang Lu, and Jorge Nocedal. Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Softw.*, 23(4):550–560, December 1997.
- [19] Jacob Lucas, Brandoch Calef, and Trent Kyono. Recovering astronomical images with deep neural network supported bispectrum processing. In *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*. 2018.
- [20] B. Jia, K. D. Pham, E. Blasch, Z. Wang, D. Shen, and G. Chen. Space object classification using deep neural networks. In *2018 IEEE Aerospace Conference*, pages 1–8, March 2018.
- [21] Richard Linares. Space object classification using deep convolutional neural networks. 07 2016.
- [22] Robert Furfaro, Richard Linares, and Vishnu Reddy. Space objects classification via light-curve measurements: Deep convolutional neural networks and model-based transfer learning. 09 2018.
- [23] Kevin Schawinski, Ce Zhang, Hantian Zhang, Lucas Fowler, and Gokula Krishnan Santhanam. Generative adversarial networks recover features in astrophysical images of galaxies beyond the deconvolution limit. *Monthly Notices of the Royal Astronomical Society: Letters*, 467(1):L110–L114, 01 2017.

DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048

- [24] Levi Fussell and Ben Moews. Forging new worlds: high-resolution synthetic galaxies with chained generative adversarial networks. *Monthly Notices of the Royal Astronomical Society*, 485(3):3203–3214, 03 2019.
- [25] Alex Hocking, James E. Geach, Yi Sun, and Neil Davey. An automatic taxonomy of galaxy morphology using unsupervised machine learning. *Monthly Notices of the Royal Astronomical Society*, 473(1):1108–1129, 09 2017.
- [26] A. Asensio Ramos, J. de la Cruz Rodriguez, and A Pastor Yabar. Real-time multiframe blind deconvolution of solar images. 06 2018.
- [27] Hossein Talebi Esfandarani and Peyman Milanfar. NIMA: neural image assessment. *CoRR*, abs/1709.05424, 2017.
- [28] Le Kang, Peng Ye, Yi Ci Li, and David S. Doermann. Convolutional neural networks for no-reference image quality assessment. *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1733–1740, 2014.
- [29] Jacob Lucas, Michael Werth, Trent Kyono, Ian McQuaid, and Justin Fletcher. Automated interpretability scoring of ground-based observations of leo objects with deep learning. In *IEEE Aerospace*. 2020.
- [30] John M. Irvine. National imagery interpretability rating scales (niirs): overview and methodology, 1997.
- [31] Ronald G. Driggers, Paul G. Cox, and Michael Kelley. National imagery interpretation rating system and the probabilities of detection, recognition, and identification. *Optical Engineering*, 36(7):1952 – 1959, 1997.
- [32] Trent Kyono, Jacob Lucas, Michael Werth, Ian McQuaid, and Justin Fletcher. Determining multi-frame blind deconvolution resolvability using deep learning. In *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*. 2019.
- [33] Trent Kyono, Jacob Lucas, Michael Werth, Brandoch Calef, Ian McQuaid, and Justin Fletcher. Machine learning for quality assessment of ground-based optical images of satellites. *Optical Engineering*, 59(5):1 – 15, 2020.
- [34] Jacob Lucas, Trent Kyono, Michael Werth, Justin Fletcher, and Ian McQuaid. Automated resolution scoring of ground-based leo observations using convolutional neural networks. In *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*. 2019.
- [35] Christian Szegedy, Sergey Ioffe, and Vincent Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *CoRR*, abs/1602.07261, 2016.
- [36] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In Yee Whye Teh and Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR.
- [37] Michael Werth, Jacob Lucas, Trent Kyono, Ian McQuaid, and Justin Fletcher. Silo: A machine learning dataset of synthetic ground-based observations of leo satellites. In *IEEE Aerospace*. 2020.
- [38] Trent Kyono, Fiona J. Gilbert, and Mihaela van der Schaar. Multi-view multi-task learning for improving autonomous mammogram diagnosis. In Finale Doshi-Velez, Jim Fackler, Ken Jung, David Kale, Rajesh Ranganath, Byron Wallace, and Jenna Wiens, editors, *Proceedings of the 4th Machine Learning for Healthcare Conference*, volume 106 of *Proceedings of Machine Learning Research*, pages 571–591, Ann Arbor, Michigan, 09–10 Aug 2019. PMLR.
- [39] Ryan Smith. Nvidia volta unveiled: Gv100 gpu and tesla v100 accelerator unveiled, May 2017.
- [40] Intel xeon processor e5-2698 v4 (50m cache, 2.20 ghz) product specifications, 2016.

DISTRIBUTION A. Approved for public release: distribution is unlimited.
Public Affairs release approval #AFRL-2020-0048