

NOISE2NOISE DENOISING OF CRISM HYPERSENSPECTRAL DATA

Robert Platt^{1,2,3}, **Rossella Arcucci**^{1,3}, **Cédric M. John**⁴

¹Department of Earth Science and Engineering, Imperial College London

²I-X, Imperial College London

³Data Science Institute, Imperial College London

⁴Digital Environment Research Institute (DERI), Queen Mary University of London

{rp1818, r.arcucci}@ic.ac.uk, {cedric.john}@qmul.ac.uk

ABSTRACT

Hyperspectral data acquired by the Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) have allowed for unparalleled mapping of the surface mineralogy of Mars. Due to sensor degradation over time, a significant portion of the recently acquired data is considered unusable. Here a new data-driven model architecture, Noise2Noise4Mars (N2N4M), is introduced to remove noise from CRISM images. Our model is self-supervised and does not require zero-noise target data, making it well suited for use in Planetary Science applications where high quality labelled data is scarce. We demonstrate its strong performance on synthetic-noise data and CRISM images, and its impact on downstream classification performance, outperforming benchmark methods on most metrics. This allows for detailed analysis for critical sites of interest on the Martian surface, including proposed lander sites.

1 INTRODUCTION

The hyperspectral data collected by the Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) on-board NASA’s Mars Reconnaissance Orbiter (MRO) mission has allowed for the highest resolution spectral mapping of the Martian surface to date. During its operation from 2006 to 2022, CRISM captured over 33,000 targeted observations and mapped 86% of the planetary surface (Seelos et al., 2023). CRISM data have allowed for detailed mapping of the surface mineralogy, and understanding of the geological evolution of the red planet (Murchie et al., 2019).

Despite the vast quantity of imagery acquired by CRISM, much of the data processing and image interpretation is accomplished manually. Significant technical expertise is required to interpret the imagery, so the reliance on manual efforts represents a large bottleneck in effective use of this data. In addition, CRISM data quality has decreased steadily since its launch, attributed to failing cryocooling systems (Carter et al., 2013). As a result, much of the more recent imagery (>2012) contains a sufficiently high noise fraction (see Fig. 1) to obscure absorption features, which restricts mineral identification and analysis (Mandon et al., 2021). This limits our understanding of key sites of interest identified more recently, for example Oxia Planum, the proposed landing site for ESA’s ExoMars rover (Mandon et al., 2021). The scarcity of labelled or ground truth data presents an additional challenge, as even older CRISM data still contains noise from atmospheric and photometric effects. Therefore, noise-free ground truth data does not exist for Mars. This highlights the need for development of unsupervised or self-supervised approaches (Fernandes et al., 2022; Pletl et al., 2023). To that end, our work introduces a new, self-supervised deep learning architecture for denoising hyperspectral signals, with no requirement for paired training data. Additionally, our model expects target data which itself is noisy, which is critical for application to Planetary Sciences where perfect data is rarely available. The code used to develop this work can be found here: <https://github.com/rob-platt/n2n4m>

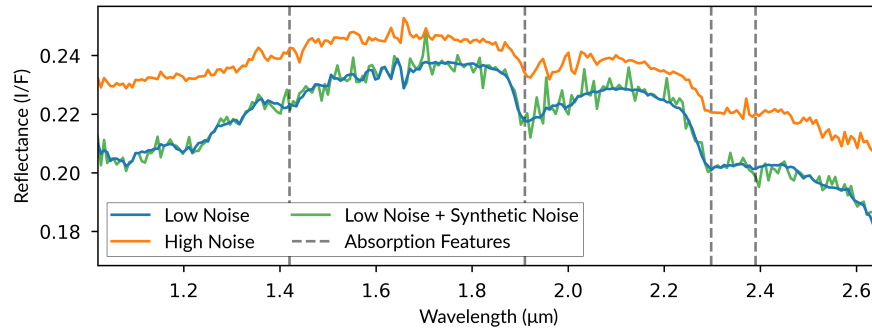


Figure 1: Comparison of a low-noise and high-noise pixel spectra of the same location extracted from two CRISM images (FRT0000A053 and FRS000364CA) acquired 7 years apart. Absorption features to identify the surface material as an Mg/Fe-Smectite are highlighted. Shown in green is the low-noise spectra with synthetic Gaussian noise added.

1.1 RELATED WORKS

Two broad approaches have been developed to reduce noise and artefacts prevalent in CRISM data. The first uses and builds upon the so-called “Volcano-Scan” correction, the current method applied in the preprocessing pipeline for CRISM IR images to correct for atmospheric effects (Murchie et al., 2022; 2009; Seelos et al., 2011). The CRISM Iterative Recognition and Removal of Unwanted Spiking (CIRRUS) and the Complement to CRISM Analysis Toolkit (CoTCAT) use statistical methods to reduce noise and build on “Volcano-Scan” products (Parente, 2008; Bultel et al., 2015). Similarly Carter et al. (2012) developed a tool (OCAT) using a Fast Fourier Transform (FFT) to generate a smoothed signal.

A second approach developed around using radiative transfer modelling, utilizing a physical forward model (DISTORT) to simultaneously post-process and denoise the spectra (McGuire et al., 2008). Developments included application of Maximum Likelihood Estimation to the results (Kreisch et al., 2017), and adding hypothesis testing for noise and data distribution (He et al., 2019). Itoh & Parente (2020) developed a method to directly estimate the atmospheric transmission from the image itself, to simultaneously apply atmospheric correction and denoising. Both groups of methods have limitations - statistical methods can over-smooth key features (Bultel et al., 2015), whilst DISTORT-based methods have higher computational requirements and long (~ 2 hrs per image) processing time (Itoh & Parente, 2020). Data-driven approaches remains unexplored.

A plethora of Machine Learning (ML) approaches have been developed for denoising of signal or image data, based on Convolutional Neural Networks (CNNs) (Tian et al., 2020; Zhang et al., 2018), Autoencoders (Gondara, 2016), and Diffusion models (Kawar et al., 2022). Of particular interest is the work of (Lehtinen et al., 2018), who implemented a CNN where the target images are not assumed to be noise free. This is critical for application to CRISM data, where there is no noise-free ground truth.

2 METHODS

The data used for this work is derived from the CRISM Machine Learning Toolkit dataset (Plebani et al., 2022). This consists of 586,599 pixels extracted from 77 low-noise CRISM images, associated with 39 corresponding mineral labels. The base Targeted Reduced Data Record (TRDR) images were processed using MarsSI (Quantin-Nataf et al., 2018) to apply standard photometric and volcano-scan atmospheric corrections. Each IR spectra contains 350 channels between $1.0210\mu\text{m}$ - $2.6483\mu\text{m}$ and $2.8070\mu\text{m}$ - $3.4769\mu\text{m}$, with channels known to contain artefacts excluded as in the original data (Plebani et al., 2022). Additionally, 150,000 spectrally inert or “bland”, pixels identified by Plebani et al. (2022) from a wider range of images were similarly processed and added to the dataset to better reflect the true class distribution. Synthetic noise is then added to the spectra to

create high-noise training examples. The model objective is to reconstruct the low-noise signal from the high-noise signal as input.

2.1 PREPROCESSING

Two further preprocessing steps were applied to clean the data - imputation of bad values ($I/F > 1$) and of data where significant atmospheric artefact remained between $1.91\mu\text{m}$ and $2.08\mu\text{m}$ due to issues with the volcano-scan correction. The data was split into train, test, and validation sets by image to ensure generalisation to unseen imagery. In addition, the Serpentine and Muscovite mineral classes were restricted to the test set only, to provide insight to the models ability to generalise to unseen mineral classes.

Synthetic noise was added to the low noise data to create training data (Fig. 1). The synthetic noise partially obscures some of the key absorption features (e.g. the absorption at $1.42\mu\text{m}$ in Fig. 1) similarly to real high noise spectra. This synthetic noise was created by sampling from a Gaussian distribution with zero mean, following Bultel et al. (2015). An additional, uniformly sampled random variable is added to increase the inter-channel variability, which better reflects the observed noise distribution (Bultel et al., 2015).

2.2 MODEL AND TRAINING SCHEME

Our approach implemented a 1-Dimensional convolutional neural network (1D-CNN) in a U-Net architecture. This was inspired by the Noise2Noise (N2N) denoising work (Lehtinen et al., 2018), which showed that when using a U-Net architecture, a model can be trained effectively to denoise natural images despite the presence of noise in the target data. Our new model architecture, Noise2Noise for Mars (N2N4M), introduces important architectural changes to fit hyperspectral data. Instead of 2-dimensional convolution, we use 1-dimensional convolution layers to leverage features in the rich spectral dimension unique to hyperspectral imagery, rather than spatial data. Transposed convolution layers were used rather than nearest-neighbour upsampling in the decoder, and kernels of size 5 were found to be the most performant during hyperparameter tuning. N2N4M was trained for 100 epochs, and the validation dataset was used to tune learning rate, batch size, and model architecture. The final model architecture and exact training parameters can be found in Appendix A. This architecture allows us to apply the model to CRISM data, where no zero-noise ground truth data exists, which has restricted data-driven approaches to the problem.

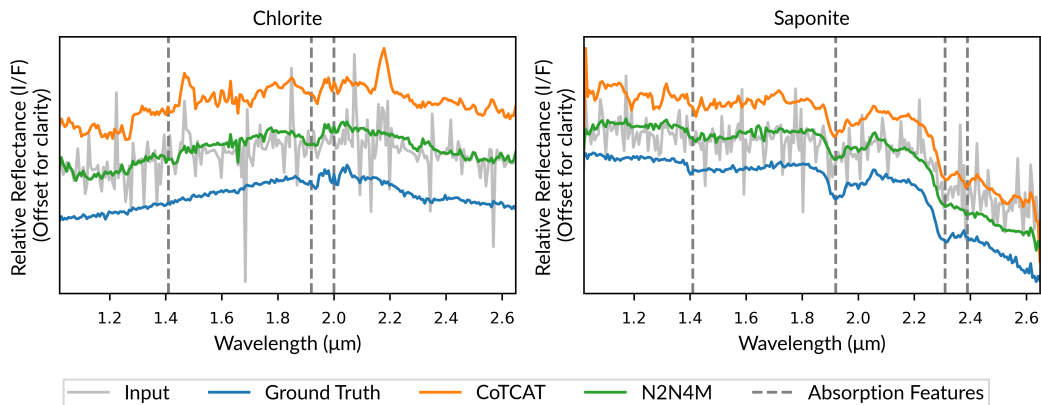


Figure 2: Qualitative comparison of results on unseen spectra from two different mineral classes. Shown are the high-noise input, low-noise target, and high-noise data denoised by CoTCAT, and N2N4M (ours). Absorption features for mineral identification are highlighted.

Data	Denoising Task	Downstream Classification Task			
	MSE	Relative Accuracy	Relative F1 Score	Relative Precision	Relative Recall
Ground Truth	N/A	1.00	1.00	1.00	1.00
Savitzky-Golay Filter	2.8×10^{-5}	0.01	0.00	0.42	0.09
CoTCAT (Bultel et al. 2015)	5.0×10^{-6}	0.35	0.43	0.50	0.41
N2N4M (Ours)	4.7×10^{-6}	0.52	0.48	0.50	0.64

Table 1: Benchmarking results for denoising, and the downstream classification task which uses the HBM from Plebani et al. (2022). Classification metrics are given relative to classifier performance on high quality, low-noise ground truth data. Best performance is highlighted in bold.

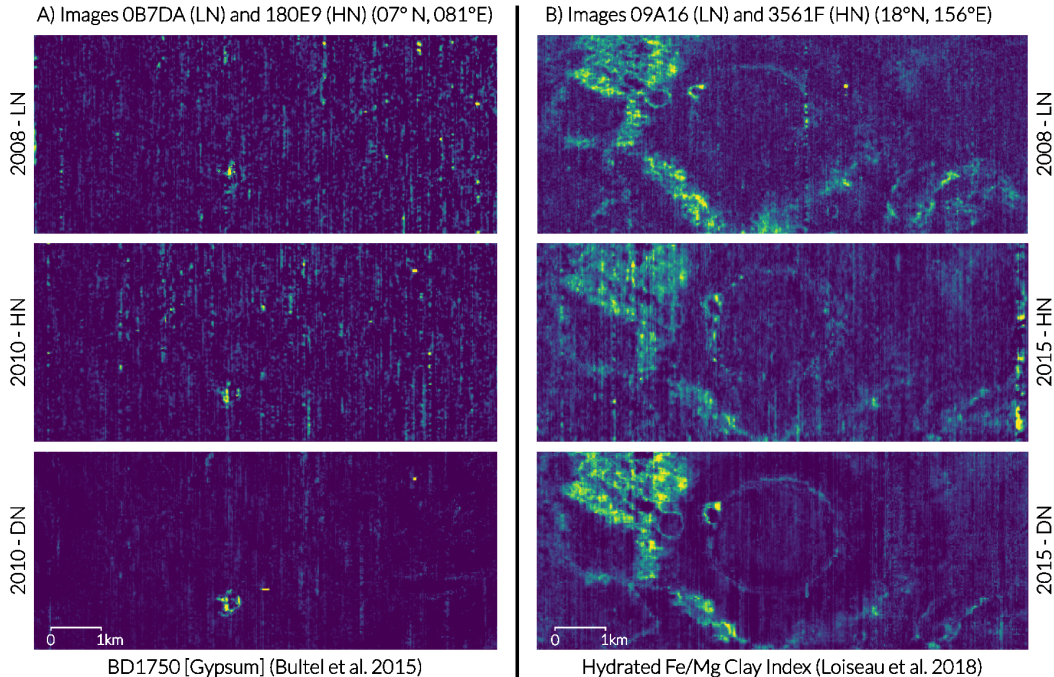


Figure 3: Two pairs of test images, from top to bottom: low-noise reference image (LN), high-noise image (HN), denoised image (DN). Images shown as the strength of absorption features (listed at base) in each pixel, which highlight outcrops of hydrated minerals identified in the reference image.

3 RESULTS

Model performance was tested with three approaches; denoising ability on synthetic noise data, downstream classification performance, and application to real imagery. The combination of these perceptual and quantitative metrics allow for a rigorous evaluation of the ability and limitations of the model. The Complement to CRISM Analysis Toolkit (CoTCAT) model (Bultel et al., 2015), and Savitzky-Golay filtering (Gorry, 1990) were used as benchmarks for comparison.

When denoising unseen synthetic high-noise data, our model shows the lowest reconstruction error compared to the ground truth (Table 1). A visual inspection of the results (see Fig. 2 and Fig. 3) shows that N2N4M removes more noise than CoTCAT whilst retaining the absorption features evident in the ground truth which are critical for mineral identification. For downstream classification, the Hierarchical Bayesian Model (HBM) developed by (Plebani et al., 2022) was used, trained on low-noise data. The HBM predicts a mineral class label for an inputted spectra, with ground truth labels included in the original dataset by (Plebani et al., 2022). The results (Table 1) are given as scores relative to performance on low-noise ground truth data, to highlight the impact of different denoising methods. N2N4M denoising results in a significant increase in most metrics over the benchmarks.

Image evaluation was performed on two CRISM image pairs. In each case, a low-noise, high quality image and high-noise, low quality image had been acquired over the same spatial location. The high-

noise image in each pair was denoised using N2N4M, before standard post-processing steps were applied to all images using the ratioing method described by Plebani et al. (2022). Summary parameters are a widespread method to visualise specific mineral suites in hyperspectral images (Viviano-Beck et al., 2014) and allow for a visual assessment of denoising quality on real data. Summary parameters for known outcrops previously identified in the high quality imagery (Bultel et al., 2015; Mandon et al., 2021) were calculated for all results (Fig. 3). The denoised images show a clear increase in perceptual quality over the original, with mineral outcrops highlighted with greater definition and clearer spatial extents. They compare favourably with the reference images and do not exhibit any spurious outcrops.

4 DISCUSSION AND FUTURE WORKS

Our proposed model shows promising results on both synthetic and real noise hyperspectral CRISM imagery. The downstream use of this data is classification of surface mineralogy, and N2N4M denoised data shows a significant improvement over the benchmark in enabling this task. As such automated classifiers are used to identify outcrop shapes and extents over exact mineral species identification (Plebani et al., 2022), the 0.23 relative increase in recall by N2N4M over the benchmark represents a large increase in capabilities. When real imagery is assessed, there is a clear increase in visual quality post-denoising. Additionally, no artificial mineral outcrops are introduced which is critical for trust in the wider use of this model outside of examples where low noise imagery is available. Whilst N2N4M shows a clear step forward in performance, further development is needed for truly reliable denoising and classification. Many higher noise CRISM images display striping noise (Carter et al., 2012), and methods to include this spatial component directly in the model should be explored.

5 CONCLUSION

A novel, data-driven method for denoising hyperspectral imagery of Mars where true no-noise ground truth is unavailable has been proposed. This approach is shown to be effective at removing noise whilst retaining key absorption features of spectra, for both synthetic noise and real example imagery. Our method is not restricted to the particular characteristics of SWIR spectra and could easily be applied in other domains e.g. Synthetic Aperture Radar (SAR), or Radio signals. Our work should allow for further analysis of much of the Martian surface mineralogy, and highlights the use-case for self-supervised ML methods in Planetary Sciences.

REFERENCES

- Benjamin Bultel, Cathy Quantin, and Loïc Lozac’h. Description of CoTCAT (Complement to CRISM Analysis Toolkit). *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 8(6):3039–3049, June 2015. ISSN 2151-1535. doi: 10.1109/JSTARS.2015.2405095.
- J. Carter, F. Poulet, S. Murchie, and J. P. Bibring. Automated processing of planetary hyperspectral datasets for the extraction of weak mineral signatures and applications to CRISM observations of hydrated silicates on Mars. *Planetary and Space Science*, 76:53–67, December 2012. ISSN 0032-0633. doi: 10.1016/j.pss.2012.11.007. URL <https://www.sciencedirect.com/science/article/pii/S0032063312003625>.
- J. Carter, F. Poulet, J.-P. Bibring, N. Mangold, and S. Murchie. Hydrous minerals on Mars as seen by the CRISM and OMEGA imaging spectrometers: Updated global view. *Journal of Geophysical Research: Planets*, 118(4):831–858, April 2013. ISSN 2169-9100. doi: 10.1029/2012JE004145. URL <https://onlinelibrary.wiley.com/doi/abs/10.1029/2012JE004145>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2012JE004145>.
- Michael Fernandes, Alexander Pletl, Nicolas Thomas, Angelo Pio Rossi, and Benedikt Elser. Generation and Optimization of Spectral Cluster Maps to Enable Data Fusion of CaSSIS and CRISM Datasets. *Remote Sensing*, 14(11):2524, January 2022. ISSN 2072-4292. doi: 10.3390/rs14112524. URL <https://www.mdpi.com/2072-4292/14/11/2524>. Number: 11 Publisher: Multidisciplinary Digital Publishing Institute.

- Lovedeep Gondara. Medical Image Denoising Using Convolutional Denoising Autoencoders. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, pp. 241–246, December 2016. doi: 10.1109/ICDMW.2016.0041. URL <https://ieeexplore.ieee.org/abstract/document/7836672>. ISSN: 2375-9259.
- Peter A. Gorry. General least-squares smoothing and differentiation by the convolution (Savitzky-Golay) method. *Analytical Chemistry*, 62(6):570–573, March 1990. ISSN 0003-2700. doi: 10.1021/ac00205a007. URL <https://doi.org/10.1021/ac00205a007>. Publisher: American Chemical Society.
- Linyun He, Joseph A. O’Sullivan, Daniel V. Politte, Kathryn E. Powell, and Raymond E. Arvidson. Quantitative Reconstruction and Denoising Method HyBER for Hyperspectral Image Data and Its Application to CRISM. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(4):1219–1230, April 2019. ISSN 2151-1535. doi: 10.1109/JSTARS.2019.2900644. URL <https://ieeexplore.ieee.org/document/8672185>.
- Yuki Itoh and Mario Parente. A new method for atmospheric correction and de-noising of CRISM hyperspectral data. *Icarus*, 354:114024, July 2020. ISSN 0019-1035. doi: 10.1016/j.icarus.2020.114024. URL <https://www.sciencedirect.com/science/article/pii/S0019103520303857>.
- Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising Diffusion Restoration Models. *Advances in Neural Information Processing Systems*, 35:23593–23606, December 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/95504595b6169131b6ed6cd72eb05616-Abstract-Conference.html.
- C. D. Kreisch, J. A. O’Sullivan, R. E. Arvidson, D. V. Politte, L. He, N. T. Stein, J. Finkel, E. A. Guinness, M. J. Wolff, and M. G. A. Lapôtre. Regularization of Mars Reconnaissance Orbiter CRISM along-track oversampled hyperspectral imaging observations of Mars. *Icarus*, 282:136–151, January 2017. ISSN 0019-1035. doi: 10.1016/j.icarus.2016.09.033. URL <https://www.sciencedirect.com/science/article/pii/S0019103516301579>.
- Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning Image Restoration without Clean Data, October 2018. URL <http://arxiv.org/abs/1803.04189>. arXiv:1803.04189 [cs, stat].
- Lucia Mandon, Adam Parkes Bowen, Cathy Quantin-Nataf, John C. Bridges, John Carter, Lu Pan, Pierre Beck, Erwin Dehouck, Matthieu Volat, Nicolas Thomas, Gabriele Cremonese, Livio Leonardo Tornabene, and Patrick Thollot. Morphological and Spectral Diversity of the Clay-Bearing Unit at the ExoMars Landing Site Oxia Planum. *Astrobiology*, 21(4):464–480, April 2021. ISSN 1531-1074. doi: 10.1089/ast.2020.2292. URL <https://www.liebertpub.com/doi/10.1089/ast.2020.2292>. Publisher: Mary Ann Liebert, Inc., publishers.
- Patrick C. McGuire, Michael J. Wolff, Michael D. Smith, Raymond E. Arvidson, Scott L. Murchie, R. Todd Clancy, Ted L. Roush, Selby C. Cull, Kim A. Lichtenberg, Sandra M. Wiseman, Robert O. Green, Terry Z. Marti, Ralph E. Milliken, Peter J. Cavender, David C. Humm, Frank P. Seelos, Kim D. Seelos, Howard W. Taylor, Bethany L. Ehlmann, John F. Mustard, Shannon M. Pelkey, Timothy N. Titus, Christopher D. Hash, and Erick R. Malaret. MRO/CRISM Retrieval of Surface Lambert Albedos for Multispectral Mapping of Mars With DISORT-Based Radiative Transfer Modeling: Phase 1—Using Historical Climatology for Temperatures, Aerosol Optical Depths, and Atmospheric Pressures. *IEEE Transactions on Geoscience and Remote Sensing*, 46(12):4020–4040, December 2008. ISSN 0196-2892. doi: 10.1109/TGRS.2008.2000631. URL <http://ieeexplore.ieee.org/document/4685531/>.
- Scott Murchie, Edward A. Guinness, and Susan Slavney. CRISM Data Product Software Interface Specification Version 1.3.7.6., June 2022.
- Scott L. Murchie, Frank P. Seelos, Christopher D. Hash, David C. Humm, Erick Malaret, J. Andrew McGovern, Teck H. Choo, Kimberly D. Seelos, Debra L. Buczkowski, M. Frank Morgan, Olivier S. Barnouin-Jha, Hari Nair, Howard W. Taylor, Gerald W. Patterson, Christopher A. Harvel, John F. Mustard, Raymond E. Arvidson, Patrick McGuire, Michael D. Smith, Michael J. Wolff, Timothy N. Titus, Jean-Pierre Bibring, and Francois Poulet.

- Compact Reconnaissance Imaging Spectrometer for Mars investigation and data set from the Mars Reconnaissance Orbiter's primary science phase. *Journal of Geophysical Research: Planets*, 114(E2), 2009. ISSN 2156-2202. doi: 10.1029/2009JE003344. URL <https://onlinelibrary.wiley.com/doi/abs/10.1029/2009JE003344>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2009JE003344>.
- Scott L. Murchie, Jean-Pierre Bibring, Raymond E. Arvidson, Janice L. Bishop, John Carter, Bethany L. Ehlmann, Yves Langevin, John F. Mustard, Francois Poulet, Lucie Riu, Kimberly D. Seelos, and Christina E. Viviano. Visible to Short-Wave Infrared Spectral Analyses of Mars from Orbit Using CRISM and OMEGA. In Janice L. Bishop, James F. Bell III, and Jeffrey E. Moersch (eds.), *Remote Compositional Analysis*, pp. 453–483. Cambridge University Press, 1 edition, November 2019. ISBN 978-1-316-88887-2 978-1-107-18620-0. doi: 10.1017/9781316888872.025. URL https://www.cambridge.org/core/product/identifier/9781316888872%23CN-bp-23/type/book_part.
- Mario Parente. A New Approach To Denoise CRISM Images. In *39th Lunar and Planetary Science Conference*, March 2008. URL <https://www.lpi.usra.edu/meetings/lpsc2008/pdf/2528.pdf>.
- Emanuele Plebani, Bethany L. Ehlmann, Ellen K. Leask, Valerie K. Fox, and M. Murat Dundar. A machine learning toolkit for CRISM image analysis. *Icarus*, 376:114849, April 2022. ISSN 00191035. doi: 10.1016/j.icarus.2021.114849. URL <https://linkinghub.elsevier.com/retrieve/pii/S0019103521004905>.
- Alexander Pletl, Michael Fernandes, Nicolas Thomas, Angelo Pio Rossi, and Benedikt Elser. Spectral Clustering of CRISM Datasets in Jezero Crater Using UMAP and k-Means. *Remote Sensing*, 15(4):939, January 2023. ISSN 2072-4292. doi: 10.3390/rs15040939. URL <https://www.mdpi.com/2072-4292/15/4/939>. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.
- C. Quantin-Nataf, L. Lozac'h, P. Thollot, D. Loizeau, B. Bultel, J. Fernando, P. Allemand, F. Dubuffet, F. Poulet, A. Ody, H. Clenet, C. Leyrat, and S. Harrisson. MarsSI: Martian surface data processing information system. *Planetary and Space Science*, 150:157–170, January 2018. ISSN 0032-0633. doi: 10.1016/j.pss.2017.09.014. URL <https://www.sciencedirect.com/science/article/pii/S0032063316304718>.
- F P Seelos, S L Murchie, D C Humm, O S Barnouin, and F Mor. Data Processing and Analysis Products Update - Calibration, Correction, and Visualisation. In *42nd Lunar and Planetary Sciences Conference*, 2011.
- Frank P. Seelos, Kimberly D. Seelos, Scott L. Murchie, M. Alexandra Matiella Novak, Christopher D. Hash, M. Frank Morgan, Raymond E. Arvidson, John Aiello, Jean-Pierre Bibring, Janice L. Bishop, John D. Boldt, Ariana R. Boyd, Debra L. Buczkowski, Patrick Y. Chen, R. Todd Clancy, Bethany L. Ehlmann, Katelyn Frizzell, Katie M. Hancock, John R. Hayes, Kevin J. Hefernan, David C. Humm, Yuki Itoh, Maggie Ju, Mark C. Kochte, Erick Malaret, J. Andrew McGovern, Patrick McGuire, Nishant L. Mehta, Eleanor L. Moreland, John F. Mustard, A. Hari Nair, Jorge I. Núñez, Joseph A. O'Sullivan, Liam L. Packer, Ryan T. Poffenbarger, Francois Poulet, Giuseppe Romeo, Andrew G. Santo, Michael D. Smith, David C. Stephens, Anthony D. Toigo, Christina E. Viviano, and Michael J. Wolff. The CRISM investigation in Mars orbit: Overview, history, and delivered data products. *Icarus*, pp. 115612, May 2023. ISSN 0019-1035. doi: 10.1016/j.icarus.2023.115612. URL <https://www.sciencedirect.com/science/article/pii/S0019103523001896>.
- Chunwei Tian, Lunke Fei, Wenxian Zheng, Yong Xu, Wangmeng Zuo, and Chia-Wen Lin. Deep learning on image denoising: An overview. *Neural Networks*, 131:251–275, November 2020. ISSN 0893-6080. doi: 10.1016/j.neunet.2020.07.025. URL <https://www.sciencedirect.com/science/article/pii/S0893608020302665>.
- Christina E. Viviano-Beck, Frank P. Seelos, Scott L. Murchie, Eliezer G. Kahn, Kimberley D. Seelos, Howard W. Taylor, Kelly Taylor, Bethany L. Ehlmann, Sandra M. Wiseman, John F.

Mustard, and M. Frank Morgan. Revised CRISM spectral parameters and summary products based on the currently detected mineral diversity on Mars. *Journal of Geophysical Research: Planets*, 119(6):1403–1431, 2014. ISSN 2169-9100. doi: 10.1002/2014JE004627. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/2014JE004627>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/2014JE004627>.

Kai Zhang, Wangmeng Zuo, and Lei Zhang. FFDNet: Toward a Fast and Flexible Solution for CNN-Based Image Denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, September 2018. ISSN 1941-0042. doi: 10.1109/TIP.2018.2839891. URL <https://ieeexplore.ieee.org/abstract/document/8365806>. Conference Name: IEEE Transactions on Image Processing.

A APPENDIX




Figure 4: Noise2Noise 4 Mars (N2N4M) Network Architecture. All convolutional kernels were size 5.

Hyperparameter settings for N2N4M model training: learning rate = 1×10^{-3} , and batch size = 256. Model was trained for 100 epochs with early stopping implemented after 5 epochs of no decrease in validation score. Mean squared error (MSE) was used as the loss function, and ADAM as the optimiser.