NOAH-H, a deep-learning, terrain analysis system: preliminary results for ExoMars Rover candidate landing sites. M.R.. Balme[1], A.M. Barrett[1], M.Woods[2], S. Karachalios[2], L. Joudrier[3], E.Sefton-Nash[3]. [1]Open University, Walton Hall, Milton Keynes, MK7 6AA, UK, matt.balme@open.ac.uk, [2]SCISYS Ltd, Methuen Park, Chippenham, SN14 0GB, UK, [3] ESTEC, European Space Agency, Noordwijk, The Netherlands.

Introduction: The ExoMars Rover [1], with a primary science goal of searching for signs of past and present life on Mars, will launch in mid-2020, and land at the Oxia Planum [2] site in early 2021. To support ExoMars Rover landing site selection, we used a deep-learning (e.g., [3]), terrain-characterisation system to analyse HiRISE [4] images, and hence produce maps of terrain types for the final two ExoMars candidate landing sites: Oxia Planum and Mawrth Valles [2]. The terrain classes were defined by the NOAH-H science team based on perceived traversability. The system was then 'trained' by providing it with 'labelled' examples of the various ontological classes. NOAH-H proved effective at distinguishing between terrain types, and was particularly good at identifying aeolian bedforms, fractured bedrock, and boulder fields.

Approach: Automated image analysis is a broad field which includes many challenging sub-tasks such as image classification, object detection and object segmentation. In NOAH-H we framed the task of terrain classification as an object segmentation problem where the goal is to classify each pixel in a HiRISE image according to a prescribed ontology defined by the science team. Deep learning [3] based neural networks provide large scale, high-dimensional, non-linear function approximation - in this case image pixels to geological class labels - derived by learning from existing labelled datasets.

To train the networks we used a tool developed as part of the parent NOAH (Novelty or Anomaly Hunter) project called the Dataset Annotation Tool (DAT [5]) to allow the science team to manually label subsets of pixel-class pairs in small 'framelets' of the larger HiRISE images. Framelets were chosen to provide representative coverage of the landing site areas. In total, DAT was used to label ~236×10^6 pixels across 1500 framelets. This dataset was used to train a variety of deep learning networks including well known backbone Convolutional Neural Network elements such as VGG16 (e.g., [6]) and ResNet101 (e.g., [7]).

The ontological classes selected were based on traversability considerations, rather than perceived geological origin or morphological landform type. The classes were informed by previous studies (e.g., [8]) but tailored for the terrains in the ExoMars candidate sites. Four main groups of terrain type were used:

(i) Bedrock. Clearly defined texture or relief suggesting outcrop. Sub-divided into 'Rugged', 'Textured', or 'Smooth' depending on the degree of metre-scale relief, and 'fractured' if containing polygonal, rectilinear or linear fractures.

(*ii) Non-bedrock*. No evidence of outcrop, interpreted to be regolith or loose materials but without bedforms or other definitive indications of aeolian origin. Sub-divided into 'Non-bedrock Smooth' or 'Non-bedrock Textured' depending on the visibility of morphological texture, and 'Non-bedrock smooth lineated', a class describing dark, mass-wasting deposits associated with scarps or impact crater rims.

(iii) Aeolian materials, defined by the presence of bedforms. Sub-divided into 'rectilinear large ripples', 'simple isolated ripples', 'simple, continuous ripples', (which describe types and spatial extents of TARs [9] which are likely to be impassable for the Rover) and 'continuous small ripples', non-continuous ripples–bedrock', and 'non-continuous ripples–non-bedrock', (which describe smaller aeolian bedforms that might or might not be traversable [10]).

(iv) Boulder fields. Identifiable by their isolated nature, relief and defined shadows.

Both landing sites had almost full coverage by HiRISE images. From these, we selected 12 (Oxia) and 18 (Mawrth) images for NOAH-H analysis, based on criteria of low-noise, full resolution (~25 cm/pixel) and central coverage of the landing ellipse. Post-processing of the NOAH-H output included downsampling to 2m/pixel and conversion to colour classes for analysis.

Results: The NOAH-H tool produced reliable, usable outputs that informed understanding of the two candidate sites, and matched closely with manual, expert-assessment mapping [2]. The results are useful for traversability studies, and also for science. NOAH-H revealed a higher proportion of bedrock in Oxia than Mawrth, revealed areas within each site where aeolian bedforms were more or less common, and was able to discriminate between continuous and non-continuous, and large and small, bedforms. It also picked out fractured bedrock areas (which correlate to some extent with clay-bearing units).

Discussion: The output for Mawrth Vallis was generally consistent from one HiRISE image to another, but in Oxia Planum there was more variability. This is partly due to the low relief in Oxia, and the resulting difficulty in differentiating between subtly different morphological classes. When similar classes were combined, the variability was almost entirely removed.

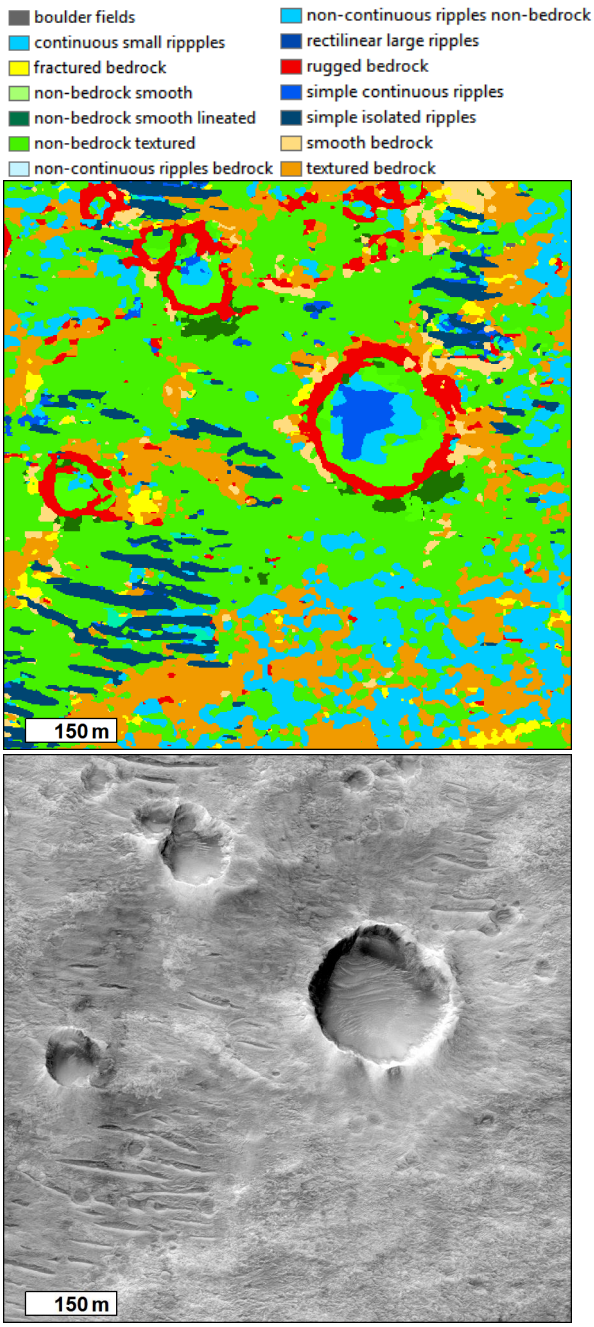

Fig. 1. NOAH-H Oxia Planum example and input HiRISE (ESP_044178_1985). Credit NASA/JPL/UoA.

NOAH-H gave the best results for morphologically distinct classes. The least well-identified classes were non-fractured 'bedrock' classes. This could be due to human difficulty in defining and/or labelling discrete and unique classes from a continuous spectrum of landscape form, and also the effects that subtle differences in lighting, noise, or look-angle can have on the ability of the NOAH-H tool to recognize different classes.

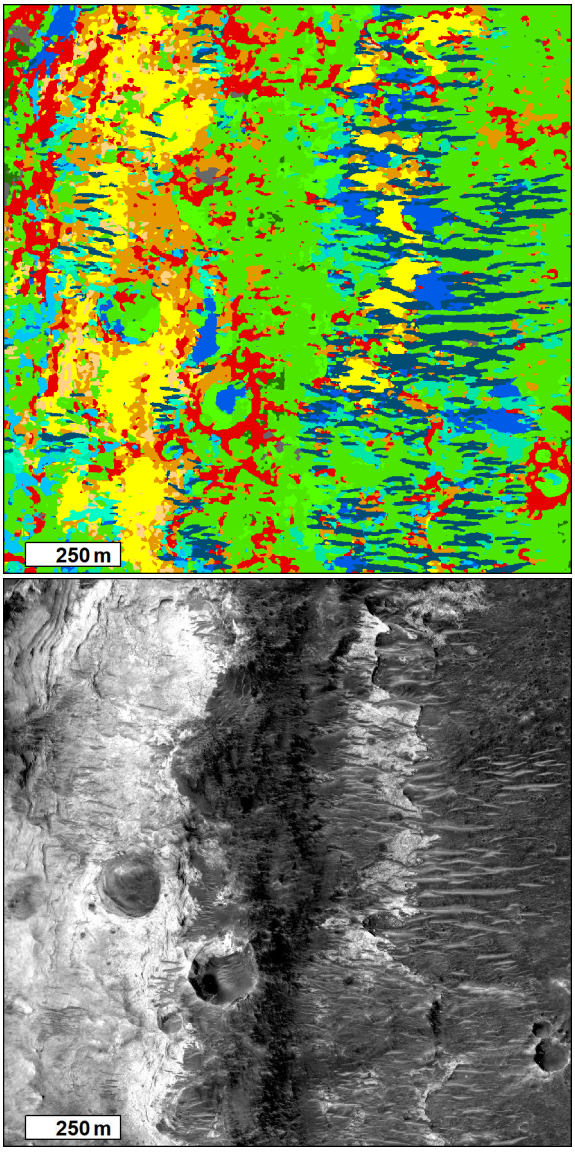

Fig. 2. NOAH-H Mawrth Vallis example and input HiRISE (ESP_046459_2025). Credit NASA/JPL/UoA.

References: [1] Vago et al., Astrobiology, 17, 6–7, pp. 471–510, 2017. [2] Loizeau et al., 50th Lunar and Planetary Science Conference, Houston, Abstr. 2378. 2019, [3] LeCun, et al., Nature, 521, 7553, 436–444, 2015. [4] McEwen et al., JGR planets, 112, E05S02, doi:10.1029/2005JE002605, 2007. [5] Wallace et al., 48th Lunar and Planetary Science Conference, Houston, Abstr. 1170, 2017. [6] Simonyan and Zisserman, arXiv:1409.1556, 2014. [7] He et al., Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778, 2016. [8] Rothrock et al., AIAA SPACE 2016, Long Beach, California, 2016. [9] Balme et al., Geomorphology 101, 4, 703–720, 2008. [10] Balme et al., Planetary and Space Science, 153, 39–53, 2017.