Exploring CaSSIS with Machine Learning – The Search for Dust Devils on Mars. V. T. Bickel[1], S. J. Conway[2], N. Thomas[3], M. Read[3], A. Valantinas[3], E. Hauber[4], P. Grindrod[5], J. Wray[6], V. Rangarajan[7], and the CaSSIS Science Team, [1]Center for Space and Habitability, University of Bern, CH (valentin.bickel@unibe.ch), [2]University of Nantes, FR, [3]University of Bern, CH, [4]German Aerospace Center, GER, [5]Natural History Museum, UK, [6]Georgia Institute of Technology, USA, [7]University of Western Ontario, CAN.

Introduction: Dust devils are atmospheric vortices made visible by dust entrained from the martian surface [1]. Their spatial and temporal distribution, shape, and velocity provide insights into Mars' atmospheric dynamics, useful for, e.g., dust, weather and climate modelling [1,2]. Yet, most of the available orbital image datasets have not been thoroughly scanned for dust devils and/or are limited to images acquired at a fixed local time (sun-synchronous orbits), limiting the representativeness of current analyses. Here, we use machine learning to systematically map dust devils in the CaSSIS (Colour and Stereo Surface Imaging System) image dataset. CaSSIS is a color and stereo imaging system onboard ESA's Trace Gas Orbiter [3] (on a non-sun-synchronous orbit) with a nominal spatial resolution of 4.6 m, covering ~6 % of Mars' surface as of November 2022 (33,130 images).

Methods: We collate all dust devil instances currently known to the CaSSIS science team (n=62). Each instance is annotated with a rectangular bounding box, referred to as labels. All labels are split into a training (n=57) and validation set (n=5) and used to tune a COCO pre-trained convolutional neural network called Yolov5x (PyTorch 1.7). In total, we train the neural net over 250 epochs (t=~10 minutes), while applying ample label augmentation, including rotation, translation, scaling, and radiometric modifications (brightness, hue, etc.). We further include negative training (i.e., non-dust devil sites, n=44) and validation images (n=6) to reduce false positives during inference (deployment). The neural net achieves a mean average precision of 0.95 in the validation set. We note that both the training and validation set are extremely small – the validation performance is therefore unlikely to be representative of the inference performance.

We deploy the best iteration of the neural net in a pre-existing processing pipeline [4] that was modified to stream and process calibrated, map-projected CaSSIS NPB composites (NIR – near-infrared, PAN - panchromatic, BLU - blue). In NPB composites, dust devils feature a bright core, an adjacent shadow, and distinct color fringes that are caused by their movement between the acquisitions of the individual color channels (Fig. 1). The processing of the entire map-projected NPB dataset took ~24 h (~1041 images per hour) using a single work station with one NVIDIA RTX 3090, running on 4 individual threads.

Results: The neural net identified a total of 255 dust devils (Fig. 1). Most of the detected dust devils were observed at L_s ~= 130° and ~270°, between 11 AM and 1 PM local time, and in Amazonis Planitia (Fig. 2 & 3). The largest dust devils were observed during local noon, while their occurrence follows the seasons (Fig. 2 & 3). A number of dust devils happen to be captured by CaSSIS stereo pairs: using the time difference between the individual acquisitions (~45 s) we derive an average dust devil velocity of ~10 m/s for sites in Arabia Terra and Aonia Terra (Fig. 3). We note that the overall ratio of dust devil/number of images scanned is one order of magnitude smaller in CaSSIS (0.01) than in CTX (0.1) as found by earlier work [2].

Discussion: The abundance of dust devils around local noon agrees well with models and lander/rover observations [5], and our velocity estimates are in line with earlier orbital measurements [6,7]. We note that the training/deployment of the neural net and the review of the results took less than 30 h, demonstrating the potential of machine learning-driven exploration of planetary image science datasets.

Future work: We will continue to deploy our detector as CaSSIS continues to build up its image archive over the coming years. In addition, we plan to develop an automated method that extracts dust devil velocity and direction using prevalent color fringes.

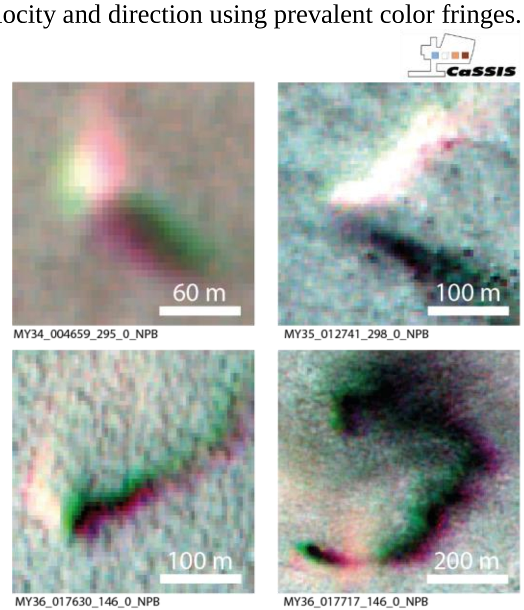

Fig. 1. Example dust devil detections (including fringes) made by the neural net in CaSSIS NPB images, North is up.

Acknowledgments: This work has been carried out within the NCCR PlanetS framework supported by SNSF grants 51NF40_182901 and 51NF40_205606. SJC acknowledges the French Space Agency, CNES.

References: [1] Reiss et al. (2012) Space Science Reviews, *203*. [2] Conway et al. (2022) LPSC 53rd, #1874. [3] Thomas et al. (2017) Space Science Reviews, *212*. [4] Bickel et al. (2020) Nature Comms., *11*. [5] Chapman et al. (2017) Icarus, *292*. [6] Stanzel et al. (2008) Icarus, *197*. [7] Reiss et al. (2014), Icarus *227*.

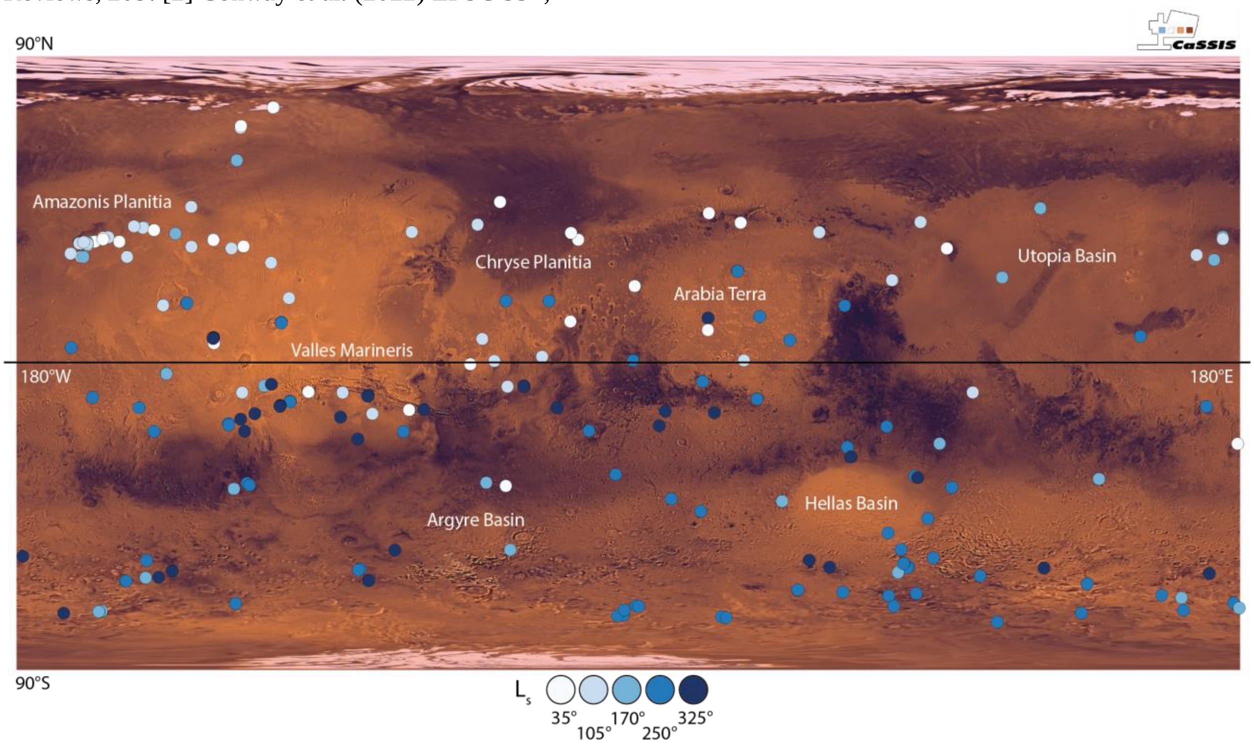

Fig. 2. Global map of neural net dust devil detections in CaSSIS. The colors indicate the L_s at which the respective dust devil was observed by CaSSIS. A globally distinct hotspot is located in Amazonis Planitia, confirming earlier findings [2]. The occurrence of dust devils appears to follow the seasons. Viking mosaic in the background.

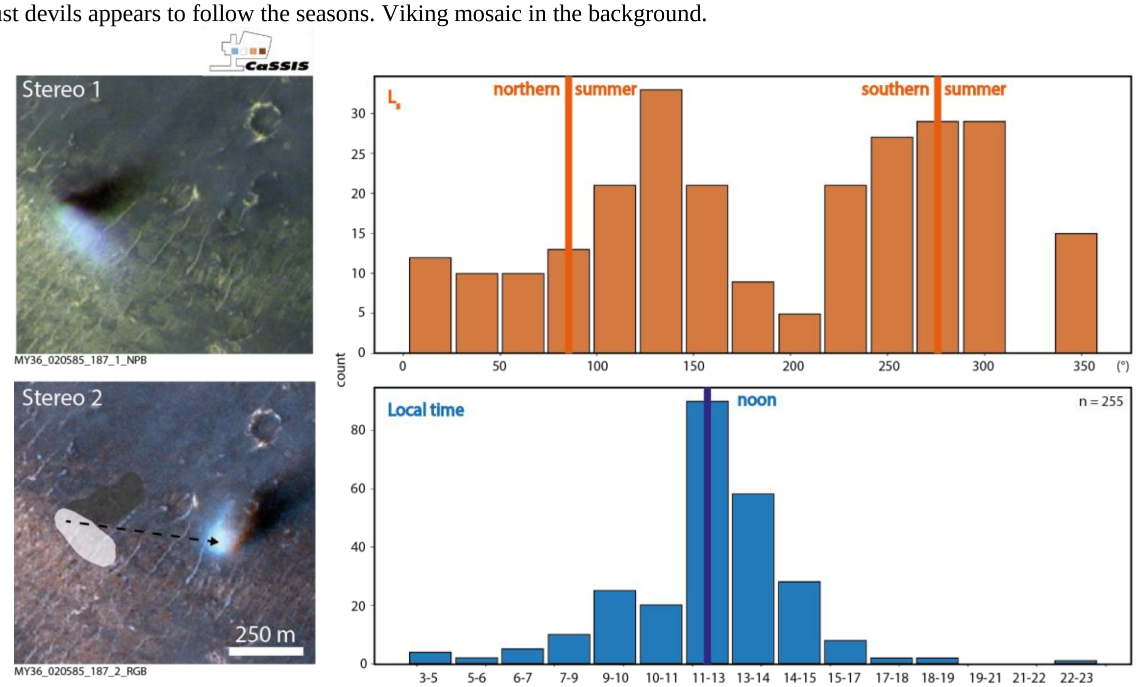

Fig. 3. Histograms of L_s (orange) and local time (blue) for the identified dust devils; CaSSIS stereo pair of a dust devil in Arabia Terra with an estimated speed of ~10 m/s (Δt ~45s). Note how its shape changes over time and depending on the viewing angle.