AUTOMATIC DETECTION OF INTERSTELLAR DUST IMPACT CRATERS ON STARDUST ALUMINUM FOILS BY CONVOLUTIONAL NEURAL NETWORKS. Logan Jaeger[1], Anna L. Butterworth[1], Zack Gainsforth[1], Robert Lettieri[1], Augusto Ardizzone, Michael Capraro, Mark Burchell[2], Penny J. Wozniakiewicz[2], Ryan C. Ogliore[3], Bradley T. De Gregorio[4], Rhonda M. Stroud[4], Andrew J. Westphal[1] , [1]Space Sciences Laboratory, University of California, Berkeley, CA 94720, [2]Centre for Astrophysics and Planetary Science, University of Kent, Canterbury CT2 7NH, UK, [3]Washington University, St. Louis, MO, 63130, [4]Materials Science and Technology Division, Code 6366, US Naval Research Laboratory, Washington, DC 20375 .

Introduction: NASA's Stardust mission successfully returned cometary and interstellar dust collectors to Earth in 2006, after an encounter with the comet 81P/Wild 2 in 2004, and exposure of the secondary collector to the interstellar dust stream in 2000 and 2002.

Aluminum foils lining these collectors also preserved impinging dust particles as impact residues within small craters, typically with diameter < 1 μm in the case of foils from the interstellar collector [1]. These craters are sparsely distributed, only detectable by scanning electron microscope (SEM) imaging, and can be easily confused with imperfections in the foil, which makes the task of locating them nontrivial.

Previous successful attempts at automatic crater identification include a normalized cross-correlation and template matching algorithm [2] and an algorithm based on a circular Hough transform and Canny edge detection [3]. Convolutional neural networks (CNNs), which have been established as effective tools in image classification [4], have the opportunity to complement or outperform these previously developed machine learning techniques. This is because, unlike the previously applied methods, CNNs can be trained on craters in combination with other features, like scratches and image exposure, and can integrate those observations into the final metric. (See Fig. 1).

Another approach to finding craters, based on engaging human volunteers to locate craters on the Stardust foils, is the website Stardust@Home (hereafter SAH) [5]. SAH has been particularly successful with identification of particle capture tracks in aerogel cells from the Stardust interstellar collector, attracting more than 30,000 citizen scientist volunteers, but a similar level of interest for investigating the interstellar foils has not emerged. CNNs may supplement manual searches like the SAH website, leading to more efficient and effective use of the labor of scientists and volunteers.

Image Creation: The images used to train the network were taken from interstellar foils I1009N and I1126N. Due to the small number of actual craters that have thus far been observed in the SEM images of these foils, this project would be impossible without analog craters produced by alternative means, namely light gas gun shots of aluminum foils. Forty four such craters had originally been imaged for use by [6], and had previously been used by [2]. We produced a further 53 analog craters with the two-stage light gas gun (LGG) in the Centre for Astrophysics and Planetary Sciences at the University of Kent using the procedure described in [7]. A training set of 20,000 crater images was generated from these "seed" craters by standard data augmentation techniques in the field of deep learning and computer vision, known collectively as data warping [8]. This process involved isolating the craters from SEM images, creating multiple versions of each crater by altering the aspect ratio, re-scaling, and rotating, and finally re-combining with a "blank" foil SEM image. Examples of images produced in this way can be seen in Fig. 2.

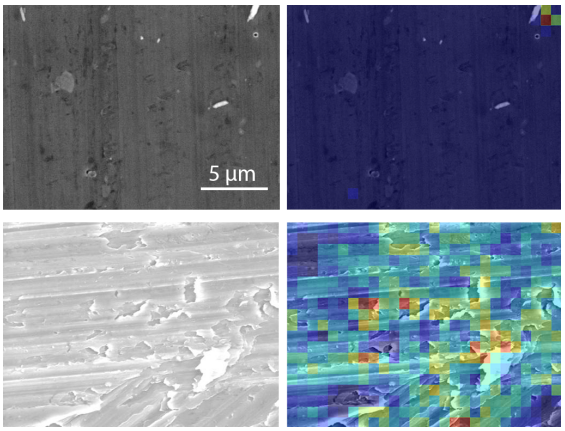

Figure 1: Each layer of the CNN contains many different convolutional filters, each with a different application. Some "light up" on crater-like features (top), others on scratches in the background (bottom). In this way, a CNN can account for the relationship between a crater-like object and its environment to better determine whether or not it truly is a crater.

Neural Network Architecture and Training: We programmed the CNN in Keras with a Tensorflow backend using Python. The CNN had a total of 216,465 trainable parameters, 8 two-dimensional convolutional layers and 7 densely connected layers. All hidden layers are immediately followed by the ReLU activation function, while the output layer is equipped with a sigmoid activation function. We used dropout and ridge-regression L2-loss to prevent overfitting.

The images in the SAH site consist of 384×512 pixel grayscale images, but it is computationally costly to train on images of that size. To alleviate this cost, we allowed our network to accept a variable-size grayscale input, and trained in multiple steps. In the first step, we trained on up to 20,000 synthetic 150×150 pixel grayscale images (10,000 each for images with and without craters). In the second step, we retrained the network, starting from the previously obtained weights, on up to 20,000 grayscale images sized 384×512 pixels. For our optimizer we used Adam with Nesterov momentum and used binary cross-entropy for our loss function.

Results and Discussion: Our network was able to identify at least 40 new crater candidates. Two of these are shown in Fig. 3. Both are from the foil I1008W. These both have yet to

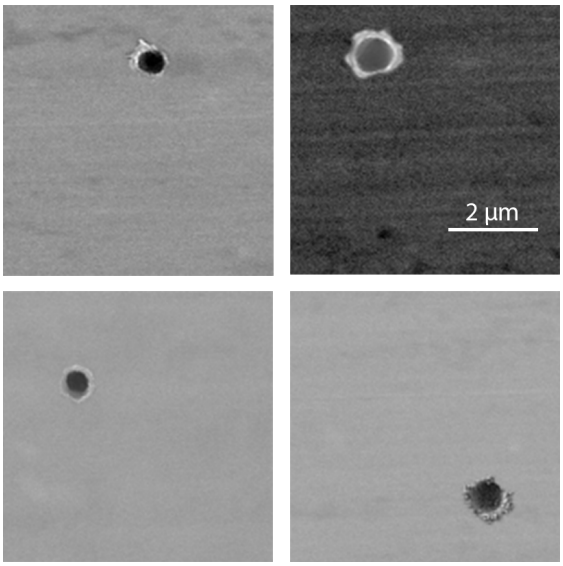

Figure 2: Examples of augmented analog craters added into images of Stardust foils that do not contain craters.

be confirmed by further SEM analysis. Neither candidate was found by [3].

The network achieved a specificity of 99.8%, measured on a sample of 500 images chosen at random from the SAH website which were confirmed not to contain craters. On a validation set of 500 images that contain the analog craters pasted into backgrounds without craters (i.e., a dataset that reflects the way in which the network was trained), the network achieved a sensitivity of 99.8%. However, in order to adequately evaluate the network, we also tested the network performance on images of "true" craters. "True" craters are those which have been determined, through high resolution scans, FIB/TEM, or any other means, to be a real impact crater. Our network correctly identified 18 of the 27 images that meet this criteria. This represents an improvement over previously developed crater-search algorithms, as the algorithm developed by [2] was able to correctly identify 8 of the 27 images. Our network also represents an improvement in speed: while the algorithm developed by [2] took about 4 minutes per processor core to search each image, our network can process 10^4 images/minute using an NVIDIA RTX 2060.

Because the network has been established as reliable, we can effectively eliminate those images that did not receive a significant prediction from the network. Less than 15,000 ($\approx 4\%$) of the roughly 424,000 images in the SAH database meet a threshold prediction value of 0.01. By removing all but those 15,000 images from the SAH website, volunteers will need to sift through significantly fewer images without craters, and their time will be utilized more efficiently.

There is a difference of approximately 30 percentage points between our network's sensitivity on analog craters and its sensitivity on true craters. This is likely due to the fact that data augmentation, though effective, cannot replace the simple act of collecting more raw data [9]. Detection sensitivity of the network could be improved by creating more analog craters with the LGG and recording more craters on each of the LGG foils, increasing the diversity of crater morphologies within the training set. Actual Stardust interstellar foil craters can also be added to the training dataset, including those detected by the current version of the network. False positives (e.g., pits) detected by the current network are also helpful, as they can be incorporated correspondingly into an improved training set to re-train the network and improve sensitivity over time.

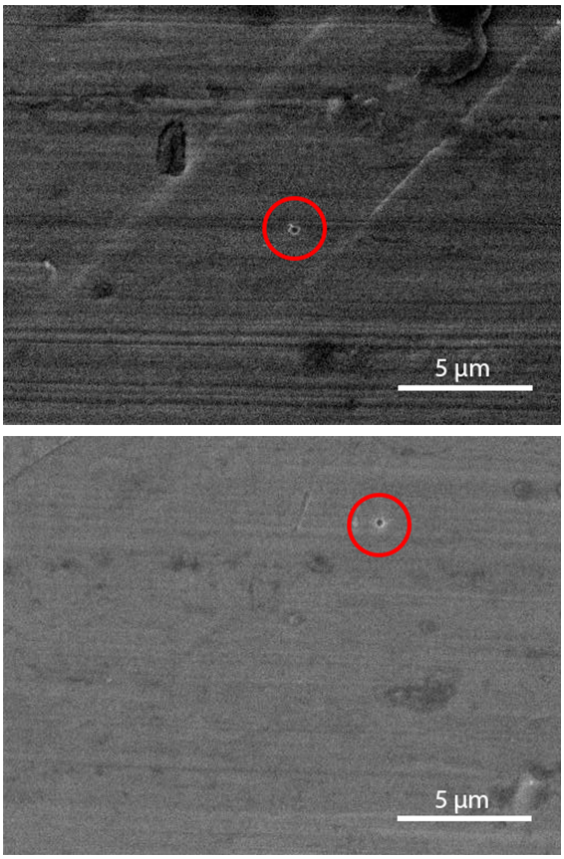

Figure 3: Examples of crater candidate images found by our network.

References: [1] Stroud R. M. et al. (2014) *M&PS*, 49, 1698-1719. [2] Ogliore R. C. et al. (2012) *M&PS*, 47, 729-736. [3] De Gregorio B. T. et al. (2018) *LPSCXLIX*, 2083. [4] He K. et al. (2015) *IEEE International Conference on Computer Vision*, 1026-1034. [5] Westphal A. J. et al. (2014) *M&PS*, 49, 1509-1521. [6] Stadermann F. J. et al. (2008) *M&PS*, 43, 299-313. [7] Wozniakiewicz P. J. et al. (2012) *M&PS*, 47, 708-728. [8] Shorten C. and Khoshgoftaar T. M. (2019) *Journal of Big Data*, 6, 60. [9] Wong S. C. et al. (2016) *IEEE* 1-6.