



Mars-DL: Demonstrating feasibility of a simulation-based training approach for autonomous Planetary science target selection

Gerhard Paar¹, Christoph Traxler², Rebecca Nowak², Filippo Garolla³, Andreas Bechtold⁴, Christian Koeberl⁵, Miguel Yuste Fernandez Alonso¹, and Oliver Sidla³

¹JOANNEUM Research, Graz, Austria

²VRVis Zentrum für Virtual Reality und Visualisierung Forschungs-GmbH, Austria

³SLR Engineering GmbH, Austria

⁴Natural History Museum, Vienna, Austria

⁵University of Vienna, Austria

Scope

Planetary robotic missions contain vision instruments for various mission-related and science tasks such as 2D and 3D mapping, geologic characterization, atmospheric investigations, or spectroscopy for exobiology. One major application for computer vision is the characterization of scientific context, and the identification of scientific targets of interest (regions, objects, phenomena) for being investigated by other scientific instruments. Due to high variability of appearance of such potentially scientific targets it requires well-adapted yet flexible techniques, one of them being Deep Learning. Machine learning and in particular Deep Learning (DL) is a technique used in computer vision to recognize content in images, categorize it and find objects of specific semantics.

In its default workflow, DL requires large sets of training data with known / manually annotated objects, regions or semantic content. Within Mars-DL (Planetary Scientific Target Detection via Deep Learning), training focuses on a simulation approach, by **virtual placement of known targets in a true context environment**.

Mars-DL Approach

The Mars-DL workflow is depicted in Figure 1, with following details:

- 3D background for the training & validation scenes is taken from 3D reconstructions using true Mars rover imagery (e.g. from the MSL Mastcam instrument).
- Scientifically interesting objects to be virtually placed in the scene are chosen to be 3D reconstructions of shatter cones (SCs) and/or Meteorites as representative set. The SCs used for 3D reconstruction are from different terrestrial impact structures and cover a wide range of lithologies, including limestone, sandstone, shale and volcanic rocks. Typical features of SCs such as striations and horsetail structure are well developed in all selected specimens and allow a clear identification. The involved stony and iron meteorites are museum quality specimens (NHM Vienna) with fusion crust and/or regmaglypts.
- A capturing campaign at the NHM Vienna covered more than a dozen SC and meteorite specimen, imaged from various positions under approximate ambient illumination conditions. This

was followed by a 3D reconstruction process using photogrammetric techniques (Figure 2) to obtain a seamless watertight high resolution mesh and albedo map for each object, gaining a 3D data base of objects to be randomly placed in the realistic scenes.

- Simulation of rover imagery for training is based on PPro3D, a viewer for the interactive exploration and geologic analysis of high-resolution planetary surface reconstructions [1]. It was adapted to efficiently render large amounts of images in batch mode controlled by a JSON file. SCs are automatically positioned on a chosen surface region with a Halton sequence, whereby collisions are avoided. The SCs are color adjusted to perfectly blend into the scene and in a future version will also receive and cast shadows. Special GPU shaders provide additional types of training images such as masks and depth maps. A typical rendering data set to be used for training is depicted in Figure 4.
- A random pose generator selects representative yet random positions and fields-of-view for the training images.
- Unsupervised generative-adversarial neural image-to-image translation techniques (GAN; [2]) are then applied on the background of the generated images, thus producing realistic landscapes drawn from the domain of actual Mars rover datasets while preserving the integrity of the high-resolution shatter cone models.
- The candidate meta-architectures for the machine learning approach are **Mask R-CNN** (instance segmentation) and **Faster R-CNN** (object detection), as these currently belong to the best performing architectures for these tasks. These meta-architectures can be configured with different backbone networks (e.g. ResNet-50, Darknet-53), which provide a tradeoff between network size and convergence rate.
- A remaining unused training set will be used for evaluation.

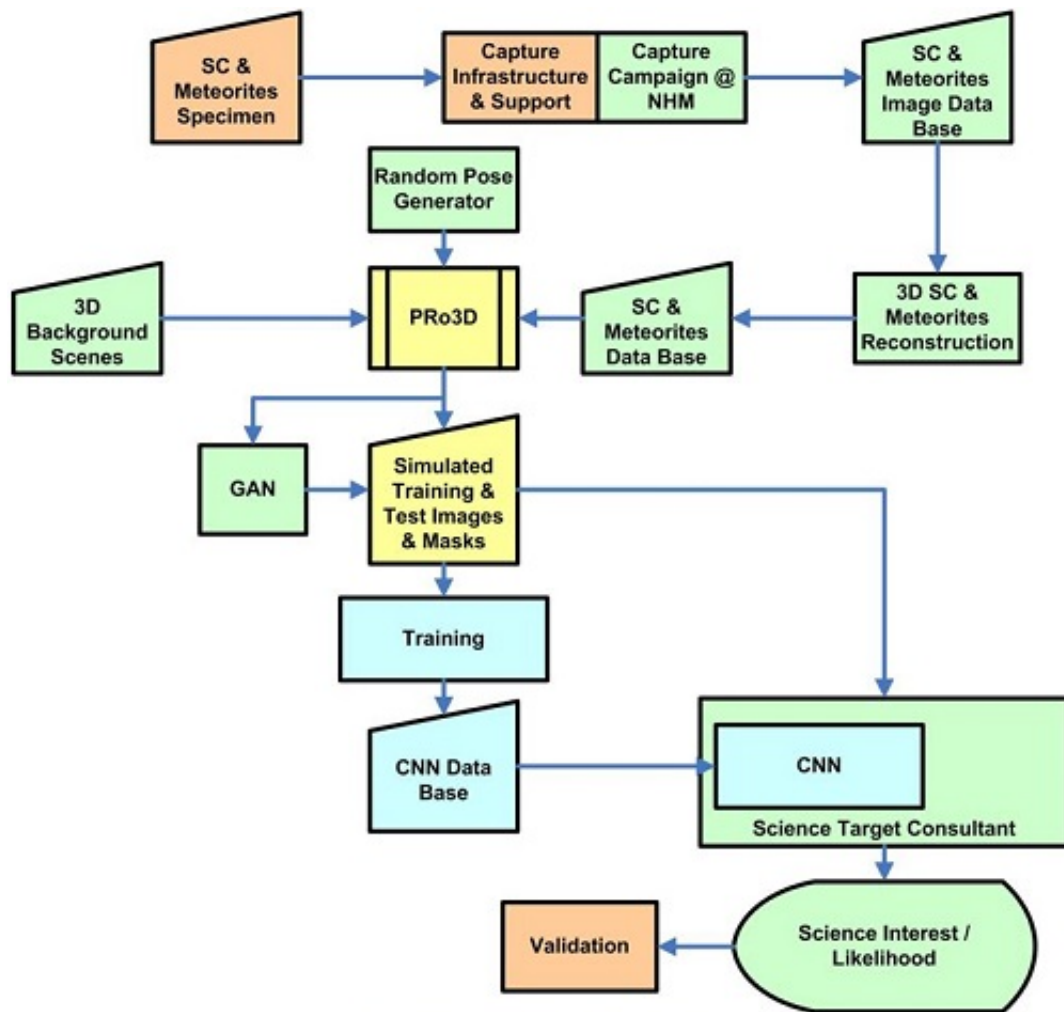


Figure 1: Mars-DL Workflow




Figure 2: 3D reconstruction of science interest target, showing capturing positions




Figure 3: Example for the impact of using a GAN to raise realism




Figure 4: Example for a simulated scene and data sets used for training. Top: Shatter cone embedded in Martian 3D reconstruction. Middle: Background scene only. Bottom: Mask of shatter cone used for training input and/or validation

Summary and Outlook

So far science autonomy has been addressed only recently with increased mobility on planetary surfaces and the upcoming need for planetary rovers to react autonomously to given science opportunities, well in view of limited data exchange resources and tight operations cycles.

An Austrian Consortium in the so-called "Mars-DL" project is targeting the adaptation and test of simulation and deep learning mechanisms for autonomous detection of scientific targets of interest during robotic planetary surface missions.

For past and present missions the project will help explore & exploit further existing millions of planetary surface images that still hide undetected science opportunities.

The exploratory project will assess the feasibility of machine-learning based support during and after missions by automatic search on planetary surface imagery to raise science gain, meet serendipitous opportunities and speed up the tactical and strategic decision making during mission planning. An automatic "Science Target Consultant" (STC) is realized in prototype form, which, as a test version, can be plugged in to ExoMars operations once the mission has landed.

In its remaining runtime until end of 2020, Mars-DL training and validation will explore the

possibility to search scientifically interesting targets across different sensors, investigate the usage of different cues such as 2D (multispectral / monochrome) and 3D, as well as spatial relationships between image data and regions thereon.

During mission operations of forthcoming missions (Mars 2020, and ExoMars & Sample Fetch Rover in particular) the STC can help avoid the missing of opportunities that may occur due to tactical time constraints preventing in-depth check of image material.

Acknowledgments

This abstract presents the results of the project Mars-DL, which received funding from the Austrian Space Applications Programme (ASAP14) financed by BMVIT, Project Nr. 873683, and project partners JOANNEUM RESEARCH, VRVis, SLR Engineering, and the National History Museum Vienna.

References

- [1] Barnes R., Sanjeev G., Traxler C., Hesina G., Ortner T., Paar G., Huber B., Juhart K., Fritz L., Nauschnegg B., Muller J.P., Tao Y. and Bauer A. Geological analysis of Martian rover-derived Digital Outcrop Models using the 3D visualisation tool, Planetary Robotics 3D Viewer - PRo3D. In Planetary Mapping: Methods, Tools for Scientific Analysis and Exploration, Volume 5, Issue 7, pp 285-307, July 2018
- [2] Jun-Yan Zhu, Taesung Park, Phillip Isola, Alexei A. Efros; Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. The IEEE International Conference on Computer Vision (ICCV), 2017, pp. 2223-2232