



Simulating rover imagery to train deep learning systems for scientific target selection

Christoph Traxler¹, Laura Fritz¹, Rebecca Nowak¹, Gerhard Paar², Christian Koeberl³, Andreas Bechtold^{3,4}, Filippo Garolla⁵, and Oliver Sidla⁵

¹VRVis Zentrum für Virtual Reality und Visualisierung Forschungs-GmbH, Geospatial Visualization, Vienna, Austria
(traxler@vrvis.at)

²JOANNEUM Research, Austria (gerhard.paar@joanneum.at)

³University of Vienna, Austria (christian.koeberl@univie.ac.at)

⁴Museum of Natural History Vienna, Austria (a.bechtold@gmx.at)

⁵SLR Engineering GmbH, Austria (os@slr-engineering.at)

1 Introduction & Scope

The remarkable success of deep learning (DL) for object and pattern recognition suggests its application to autonomic target selection of future Martian rover missions, to support planetary scientists and also exploring robots in preselecting possibly interesting regions in imagery, increase the overall scientific discoveries, and to speed-up the strategic decision-making.

Deep learning requires large amounts of training data to work reliably. Many different geologic features are to be detected and to be trained for in a DL-systems. Past and ongoing missions such as the Mars Science Laboratory (MSL) do neither provide the necessary volume of training data nor existing “ground truth”. Therefore, realistic simulations are required. The scheme and workflow of the Mars-DL simulation is depicted in Figure 1.

Realistic simulations are based on accurate 3D reconstructions of the Martian surface, which are obtained by the photogrammetric processing pipeline PRoVIP [1]. The resulting 3D terrain models can be virtually viewed from different angles and hence can be used to obtain large volumes of training data.




Figure 1: Simulation workflow diagram.

2 Shatter Cones

For Mars-DL it was decided to focus on shatter cones (SC), which are macroscopic evidence of shock metamorphism and form during meteorite impact events [2]. Shatter cones take on a variety of sizes and shapes depending mainly on the rock type, but all of them have distinctive fan-shaped “horsetail” structures that resemble striations, which makes them well-suitable objects to train a DL-system and assess its detection reliability. Due to their rare presence in real imagery, we place textured 3D models of SCs into the real context of a Martian surface region. A 3D capturing and reconstruction campaign provided a series of 3D models from SCs taken from actual terrestrial impacts (Figure 2).

They have sufficient geometric and texture resolutions to preserve characteristic features. To allow shading of SCs the texture was captured with homogeneous lighting to obtain a pure albedo map without any prebaked lighting. High resolution DTMs (Digital Terrain Models) of Martian surfaces usually have an image texture being derived from rover instrument imagery. Their prebaked lighting effects such as shading, shadows and specular highlights are determined by the sun’s direction at capturing time, obtained from SPICE. For a shatter cone to perfectly blend into the scene, it first needs to be shaded from the same illumination direction and second it needs to cast a shadow from the same angle as found in the original background surface.

Many shatter cones from terrestrial sites have a contact surface, where they were broken or cut from the outcrop and museum specimens often bear a label. This surface must not be visible in training images (although broken cones exist in impact deposits as well), and are marked by a vector as part of the meta-data of the 3D model to keep them invisible during automatic positioning.





Figure 2: Capturing shatter cones with homogeneous lighting.

3 Simulation with PRo3D

Planetary scientists can explore high-resolution DTMs with a dedicated viewer called PRo3D [3]. Fluent navigation through a detailed geospatial context allows visual experience close to field investigations. Therefore, PRo3D offers much of the required functionality to generate large volumes of training images.

PRo3D was extended for batch rendering to enable mass-production of training data. It accepts commands from a JSON file that defines many different viewpoints from which to render different types of images. The training of DL systems also demands masks for the shatter cones and depth images. For both, special GPU shaders were implemented.

The automatic random positioning of shatter cones is also controlled by JSON commands. We calculate positions for a given viewpoint region with a Halton sequence, which guarantees an even distribution. The size and rotation of each SC is randomly modified within given ranges, observing visibility constraints mentioned above. Axis-aligned bounding boxes of SCs are used to determine how deep they penetrate into the surface and allow efficient collision detection between SCs. Intersections are avoided by removal of individual SCs (Figure 3). To enhance realism, SCs are color adjusted to the image texture of the surrounding environment and shaded with the appropriate illumination direction obtained from the corresponding SPICE kernels (Figure 4).



4 Conclusion

We present an efficient way to generate large amounts of images for the training of DL systems to support autonomous scientific target selection in future rover missions. It is based on available high-resolution Martian surface reconstructions and newly created shatter cone models. The viewer PRo3D, developed for planetary science explorations, was extended for the efficient mass simulation of different DL training sets.

Future work includes the design of a shader that combines shadows cast from SCs with the image texture of the surrounding surface and a fully automatic color adjustment.

References

[1] Paar G., Muller J.P., Tao Y., Pajdla T., Giordano M., Tasdelen E., Karachevtseva I., Traxler C., Hesina G., Tyler L., Barnes R., Gupta S., Willner K. PProViDE: Planetary Robotics Vision Data Processing and Fusion. In EPSC Abstracts, Vol. 10, EPSC2015-345, European Planetary Science Congress 2015.

[2] French B.M., and Koeberl C. (2010) The convincing identification of terrestrial meteorite impact structures: What works, what doesn't, and why. *Earth-Science Reviews* 98, 123–170.

[3] 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.

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.