

EXPLORING AND EXPLOITING LARGE FIELD TRIAL DATASETS FOR PERCEPTION AND SIMULATION

Iain Wallace⁽¹⁾, Mark Woods⁽¹⁾

⁽¹⁾ SCISYS, 23 Clothier Road, Bristol, BS4 5SS, UK, {firstname}.{lastname}@scisys.co.uk

ABSTRACT

The collection of large scale, detailed, data products during field trial or live rover operations opens a wealth of possibilities in autonomy, simulation and the potential to support future innovation. Processing and exploiting these data has been a common theme through our work in recent years which is explored and presented in this paper.

Field datasets are essential for the development of high-quality science autonomy, high fidelity simulation and system testing. Finally, their automated analysis presents tantalising possibilities for assisting in the dissemination and analysis of science data returned from exploration missions.

1. INTRODUCTION

SCISYS have run and been partners in many ESA, UKSA and other projects that have involved a field trial to some degree. These range from weeks spent in the Atacama Desert Martian analogue for SEEKER[1], SAFER[2] and Chameleon[3] to local UK trials such as PMOPS [4] and spin-out applications of space robotic technology like DEEPRECALL – identifying defects in tunnels. The automated, robotic, capture of image sequences, lidar and other sensor data produces enormous data, which has utility beyond the original capture purpose. To give some indication of volume, approximately 10 days of trials in Chameleon gave rise to ~1 Tb of data products. Furthermore, these datasets may be combined and enhanced in a variety of ways

with some post-processing.

One illustrative large dataset is the SCISYS AtacamaOneMillion image dataset. This is a set of ~1.1 million images from the left camera of a stereo navcam pair collected over the SEEKER[1], SAFER[2] and Chameleon[3] Atacama field trials. This large data can be used for many testing and evaluation purposes, as well as experimentation into cutting-edge machine learning approaches. For example, Figure 1 shows the output from a Deep Convolutional Generative Adversarial Network[5] that has learnt a good representation of the desert imagery from unsupervised input.

Enabling data-driven experiments such as these into deep learning will be the future of machine perception for space exploration. The recent explosion in the capabilities and applications of deep convolutional neural networks (CNNs) for machine vision makes it impossible to ignore their promise for the many vision challenges for autonomy in space exploration. The ESA NOAH project, which follows on our earlier MASTER[6] work, will experiment with the application of CNNs and the challenges with achieving flight implementations.

We have already had some success applying CNNs to these data for the purposes of allowing users to browse and understand large datasets in the UKSA funded Deep Recall project. We present an example of how they can be used to extract semantic understanding of the image

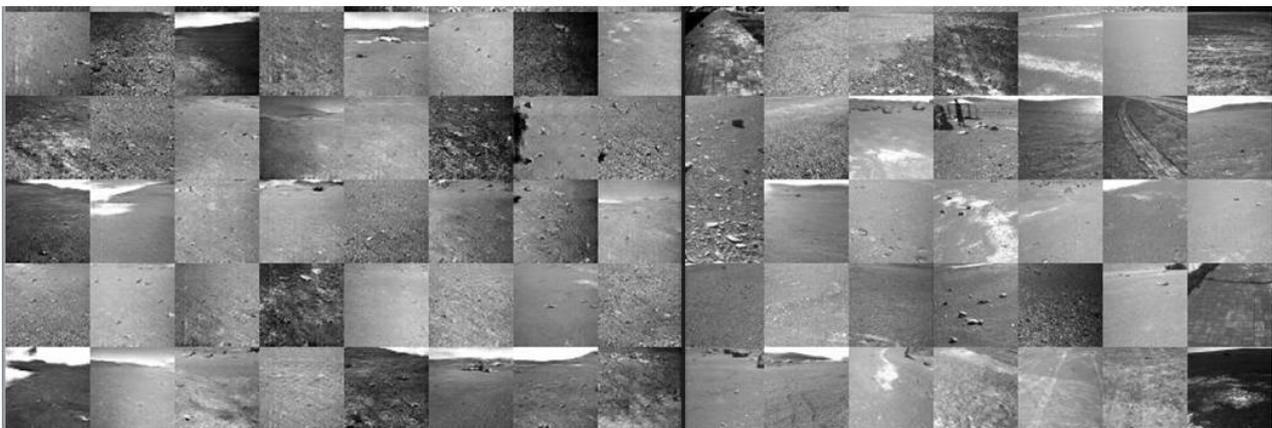


Figure 1 - Examples of synthetic terrain images generated by a DCGAN (left) and training images from the AtacamaOneMillion dataset (right). The realism of the generated images illustrates that the network can learn an appropriate feature space.

content and allow easy browsing of the AtacamaOneMillion data.

Collection of high quality field data also allows for improved fidelity simulation for experimentation, verification and validation. With the example of the SCISYS ESIM software, used on recent ESA and UKSA projects, we illustrate how these datasets can also be used to facilitate interactive autonomy simulation and experimentation.

2. DATASETS

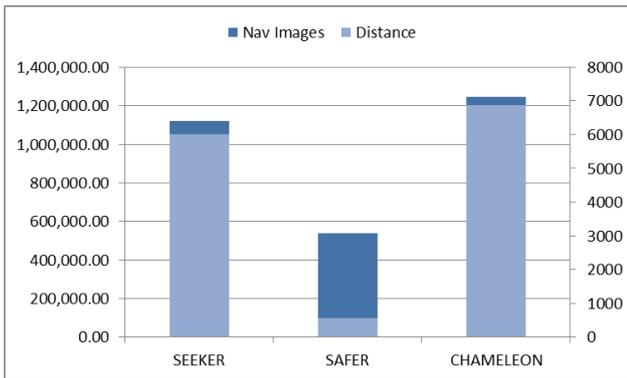


Figure 2 - Number of Navcam Images From Field Trials

Figure 2 gives some indication of the volume of image data that might be acquired on a trial. The number of Navcam images is plotted with the distance of autonomous traverse. SAFER was focussed on science operations, whereas the other trials focussed on autonomous navigation with faster rovers – hence the difference in ratio of distance to number of images. And these image data are just a portion of the overall – Chameleon also collected 2,300 UAV photographs, 225,000 point clouds plus a large volume of rover platform telemetry and ground truth DGPS data and so on.

A key limitation in exploiting datasets is the lack of metadata for future analysis. It is normal for significant data to be recorded alongside images for example, but if it later becomes useful process the data according to different criteria for unforeseen uses then the metadata required may not be present – it is hard to imagine all possible future use. To tackle this, we can exploit our work on autonomous science and image processing, and extract semantic meaning from the content of the images – search for images with certain types of rocks, even although this metadata is not recorded. These types of applications are discussed in the next section.

To support experiments such as these we have created new test datasets, such as AtacamaOneMillion. This is a

~1.1 million monocular image dataset made up of left navigation camera images from all three of the above trial datasets. Test or engineering trial runs which don't primarily feature desert terrain have been excluded, though some images remain with non-martian like features such as people, camps or vehicles. The vast majority of the dataset is the empty Mars-like expanse of the desert. This dataset has been used to develop and test machine learning approaches for automated image understanding – for example the MASTER project uses a subset from these as one of its datasets, and the images in Figure 1 were made possible by the large volume of data supported an unsupervised-learning approach.

Datasets can also support simulation, as well as experimentation. In each Atacama trial there have been Unmanned Aerial Vehicles (UAVs) flown to collect overhead data to produce maps similar to those available using HiRISE on Mars. In addition, for Chameleon SCISYS used a multirotor platform capable of more controlled, directed flight than the previous fixed-wing UAVs to enable capture of very high resolution data. From these data products, and GPS ground truth, we can process them photogrammetrically to produce highly detailed, textured 3D models. An example 3D reconstruction is shown in shown in Figure 3 – the high-resolution source imaged enabled a reconstruction to a level of detail such that boot tread pattern is clear in footprints.



Figure 3 - View of High-Resolution 3D Terrain Model. The terrain is modelled to 5mm 3D resolution, with more detailed texture mapping.

3. DEEP LEARNING FOR SEMANTIC DATASET EXPLORATION

Recent advances in machine learning, particularly as applied to computer vision, have been made in the field of deep learning. In this section we describe recent work SCISYS have completed to apply this technology to new application domains and large datasets.

3.1. The DEEPPRECALL Project

We have also undertaken a project for the UK Space Agency called Deep Recall. This is a non-space spin out project to apply similar algorithms to our autonomous science work in an industrial inspection application. It involved the creation of a back end server processing imagery with Deep Convolutional Neural Networks and then a front end web interface by which an end user could browse and filter object detections from the imagery, generating reports, and providing feedback to filter detections and therefore improve the training data

The deployment diagram in Figure 4 illustrates the workflow of the Deep Recall System. An expert annotates examples images that are used to train a model. In a separate process the model is deployed on survey data leading to detections presented to the expert for review, and statistics or reports generated. The expert can also visually review the detections and use this as feedback into the training process, bootstrapping the system where there is a scarcity of labelled data

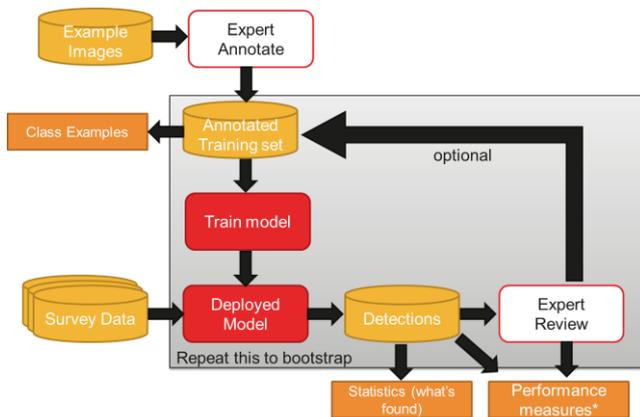


Figure 4 - Deep Recall Deployment Architecture

In the screenshot of the interface in Figure 5, a user adjusts the confidence threshold in order to browse the results of bolt pockets detected in tunnel imagery using a DCNN.

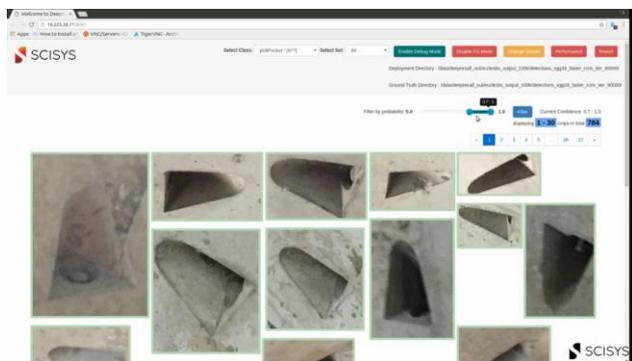


Figure 5 - Deep Recall Web Interface

Inspection data is not the only data we have processed with this system. We have also used the training data produced for the MASTER project to produce models suitable to apply to the AtacamaOneMillion. This provides an example of how meaningful semantic information can be extracted from the image content of large datasets and used to browse it.

The promise of such approaches is metadata-less browsing and processing of data products. To give a concrete example, perhaps we wish to select example navcam images where there is glare from the sun to test new stereo processing algorithms. Previously, this would require recording metadata to manually tag images, or perhaps inferring from recorded rover pose, time of day, latitude and camera angle when glare might be expected, then manually selecting from those sequences. Using our approach, we simply extract all images where glare is present, automatically. This is inferred from the pixels themselves, and so we can answer a query like this that was not anticipated at the time of data capture. Figure 6 shows the output from such a query.

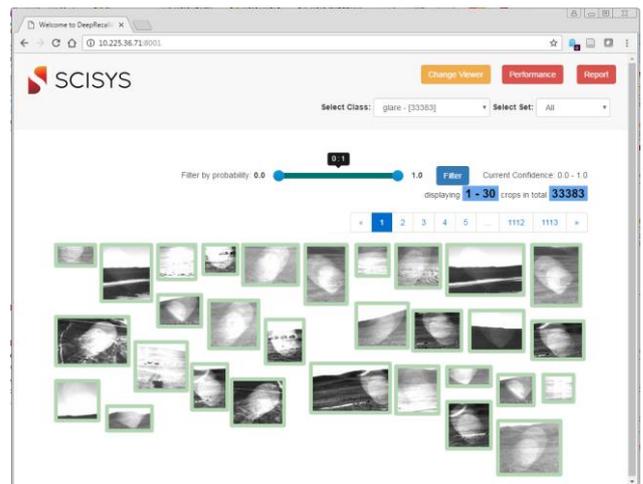


Figure 6 - Selecting only images with camera glare.

3.2. The NOAH Project

The ESA NOAH project follows on from the prior MASTER[6] project in working towards a robust autonomous science capability for European robotic exploration missions, and ground based data processing. Where MASTER was directed toward both rover and Earth Observation (EO) imagery, NOAH is focussed on Martian Rover Navcam images.

By adding autonomous capabilities to detect novel or scientifically interesting phenomena in images we can enhance the scientific return of robotic exploration missions. The many images acquired for navigation during traverse operations that are normally discarded

can be inspected, and those scoring highly selected for downlink or even further sensing of the target. Of course more powerful rover AI and autonomy will not replace ground-based scientists, but it offers the potential to lower the probability of missing proximal science targets.

The dataset driving this work is not field test data as such; rather it is a large set of operational Martian rover data from the MER and MSL rovers. We have collated a set of sixty-thousand navcam images. To process these we can filter out images primarily of engineering interest of the rover deck, or images primarily of the sky, by extracting and processing the PTU metadata. This leaves us with a very large set of Martian science imagery to experiment with. From this set, five-thousand have been selected in consultation with a planetary geologist and will be annotated with classes of interest in the online labelmars.net effort[7]. This will produce one of the largest semantically annotated image datasets of Martian Navcam imagery and serve as the fuel that drives cutting-edge autonomous science for European rovers.

As well as the activity to produce this large annotated dataset, the NOAH project also advances the TRL of our science autonomy solution. Through a program of algorithmic research and development, combined with implementation of a Prototype Flight Detector we will advance the autonomous scientist concept toward the ultimate goal of a flight experiment on ExoMars or subsequent rover missions.

As part of the algorithmic improvement effort we will address modern deep learning approaches – these represent the state of the art of computer vision systems, approaching or exceeding human performance in some tasks. Deep neural networks are attractive in that they form high-capacity, even universal, function approximators which can be queried in constant-time during inference. That is to say that the computation required is only a function of the network architecture. This is in stark contrast to other approaches with similar representational power such as non-parametric methods which carry a computational burden during inference which is a (often non-linear) function of the number of training data considered. The goal of constant-time, state of the art vision performance is certainly attractive for a resource-constrained highly dependable system – such as a planetary rover.

The combination of these new techniques and the large dataset to test on gives tantalising possibilities. There is promise not just of on-board systems to assess images but also ground systems to process and browse large image collections according to their scientific content.

4. DATASETS FOR SIMULATION AND TESTING

Whilst these large field datasets can be used for implementation and testing of various vision (or other modality) algorithms, there are limitations. If processing an image sequence as input to a complete autonomous system it cannot make different decisions to the capture platform – the traverse is the same every time. For a lot of work it is useful to be able to perform end-to-end system tests, and closes the loop from camera input to rover navigation.

To this end, SCISYS developed our ESIM tool, initially as part of the ESA HRAF project and then further extended under the UKSA HATS activity. This tool provides a high-fidelity visual simulation of a field trial environment in which a virtual rover can be driven. Figure 3 shows a user view of the ESIM, and Figure 7 provides an example of synthetic views from the ESIM.

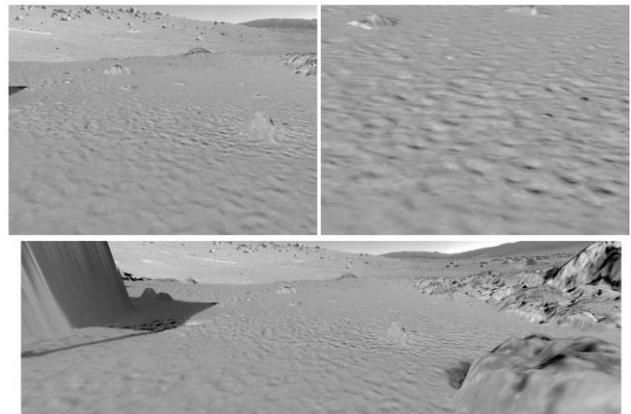


Figure 7 - Example images from a simulated normal, telephoto and wide-angle camera.

Whilst these are synthetic camera views, a key innovation is that they are based on real data acquired during field trials. In the case of Figure 3 this is Atacama trial data from the UKSA Chameleon trials, Figure 7 shows data from the Proviscout[8] trial in Tenerife. The source data is high resolution monocular camera images from either fixed-wing or multirotor UAVs. These data are then processed into detailed digital elevation maps (DEMs) using photogrammetric techniques, then stitched and orthorectified to provide textures to drape. Different camera views are defined by the intrinsic parameters recovered from the calibration processes applied to real cameras. Thus the ESIM allows for “virtual field trials” and end-to-end testing of improvements to complete autonomous systems.

5. APPLICATIONS AND FUTURE WORK

There are two main avenues of future work to consider here:

- What are the applications of the large, varied datasets SCISYS have collected?
- What are the future applications of the algorithms and software we have developed?

Access to large datasets and deep understanding of their characteristics is an absolutely vital requirement of developing autonomous systems. This has always been true, but recent advances in data-driven machine learning makes this even more the case. In this paper we have presented a few of the space-science related datasets we have assembled over the years, in addition to other space and non-space datasets such as tunnel inspection data. As well as improving algorithms for autonomous science, these datasets can themselves yield important scientific results through their analysis. For example, questions such as frequency of float rocks, or the abundance of light-toned veins can be treated statistically over a large set of images. Environmental and formation parameters can thus be characterized, e.g., the intensity of fracturing at different rover positions/landing sites. The study of the annotated Martian images from the NOAH project is expected to yield such results.

Further advances in simulation are also possible. High quality synthetic images such as those produced by GANs illustrated in the introduction hold some promise, as do further advances in the photogrammetric processing of field data to produce detailed 3D scenes.

Future developments of our autonomous science software are expected in several directions. In the near term, the NOAH project will expand capability and raise TRL. Going forward we would hope to demonstrate on flight-like hardware to progress toward an ultimate goal of a flight experiment.

In terms of applying deep learning approaches to browse large datasets, as part of this work we have developed new web based interfaces to interact with and browse this data. It is easy to see how this distributed, metadata-less browsing of data according to its content could have uses, in both space and non-space. There are applications such as browsing and disseminating data products from exploration missions to processing robotic inspection data from the built environment on Earth.

6. SUMMARY

In this paper we have discussed the vast volumes of data produced by field trial efforts, using the example of our one-million plus Atacama dataset as an exemplar. However, far from being a problem, these large datasets present an advantage, enabling cutting edge machine learning techniques for image understanding and providing high quality data for system simulation and

testing. Finally, we describe how the application of the image processing techniques these data enable allows for new interfaces to the data – browsing datasets by the semantic content of the data. This metadata-less processing will only become more important, and yield greater benefits in the future.

7. REFERENCES

- [1] M. Woods *et al.*, ‘Seeker—Autonomous Long-range Rover Navigation for Remote Exploration’, *J. Field Robot.*, vol. 31, no. 6, Oct. 2014.
- [2] M. Woods, A. Shaw, I. Wallace, Malinowski, Mateusz, and Rendell, Phillip, ‘Simulating Remote Mars Rover Operations in the Atacama Desert for Future ESA Missions’, presented at the 13th International Conference on Space Operations, 2014.
- [3] M. Woods, A. Shaw, and I. Wallace, ‘THE CHAMELEON FIELD TRIAL: TOWARD EFFICIENT, TERRAIN SENSITIVE NAVIGATION’, presented at the ASTRA 2015, 2015.
- [4] P. Rendell, I. Wallace, M. Woods, and D. Long, ‘PMOPS: The Planetary Mission On-Board Planner and Scheduler’, in *Proc. 14th Symposium on Advanced Space Technologies in Robotics and Automation*, 2017.
- [5] A. Radford, L. Metz, and S. Chintala, ‘Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks’, *CoRR*, vol. abs/1511.06434, 2015.
- [6] I. Wallace and M. Woods, ‘MASTER: A Mobile Autonomous Scientist For Terrestrial and Extra-Terrestrial Research’, *13th Symp. Adv. Space Technol. Robot. Autom. ASTRA*, no. 3, 2015.
- [7] I. Wallace, M. Woods, and N. Read, ‘LABELMARS.NET: DRIVING NEXT-GENERATION SCIENCE AUTONOMY WITH LARGE HIGH QUALITY DATASET COLLECTION’, in *Proc. 14th Symposium on Advanced Space Technologies in Robotics and Automation*, 2017.
- [8] G. Paar *et al.*, ‘PRoViScout: a planetary scouting rover demonstrator’, *Proc SPIE 8301 Intell. Robots Comput. Vis. XXIX Algorithms Tech.*, p. 83010A–83010A–14, 2012.