

5th IAA Planetary Defense Conference – PDC 2017
15-19 May 2017, Tokyo, Japan

IAA-PDC-17-09-02

APPLICATION OF MACHINE LEARNING FOR PLANETARY DEFENSE – Three Case Studies

**James Parr⁽¹⁾, Franck Marchis⁽²⁾, Michael Busch⁽²⁾, Peter Jenniskens⁽²⁾,
JL Galache⁽¹⁾, Eric Dahlstrom⁽¹⁾**

⁽¹⁾NASA FDL, 189 N Bernardo Ave #200, Mountain View, CA 94043, US,
james@frontierdevelopmentlab.org

⁽²⁾ SETI Institute, 189 N Bernardo Ave #200, Mountain View, CA 94043, US, fmarchis@seti.org

Keywords: *Machine Learning, Shape Modeling, Asteroid Deflection, Meteorite Recovery, Education*

ABSTRACT

The emergence of GPU enabled machine learning has unlocked a range of opportunities for the development of new tools and approaches for Planetary Defense. NASA's Frontier Development Lab (FDL) was established to explore the possibilities offered by the application of convolutional neural nets, Bayesian, deep generative models, machine vision and other Artificial Intelligence techniques to asteroid shape modeling, deflection efficacy and meteorite recovery. We present three projects conducted at the FDL in 2016 an applied research accelerator designed to enhance NASA's capability set by matching emerging talent from academia with peers and technologies within the private sector.

I. The NASA's Frontier Development Lab

The NASA FDL is an applied research accelerator based at NASA Ames Research Center and the SETI Institute, established to maximize the opportunity offered by proximity to technologies and capacities emerging in academia and the private sector - particularly Silicon Valley. The goal is to demonstrate how breakthroughs can be industrialized over an accelerated timeframe, in a way that is useful for America's Space Program - and, in the process, help protect the planet - specifically, from 'PHAs' or Potentially Hazardous Asteroids.



Figure 1: Logo of the NASA Frontier Development Lab, an applied research accelerator hosted at the SETI Institute.

FDL started life as a response to the Obama White House’s Asteroid Grand Challenge to NASA [1], ‘To find all asteroid threats to human populations and know what to do about them’ through unconventional approaches and partnerships - particularly the private sector. In this way, FDL was an experiment in public / private partnership from its inception and acknowledges that while NASA has specialties in the problem domain (and the data), the solution space is often a product of a collision of that knowledge with the private sector and academia - particularly in fast-moving fields like machine intelligence.

The FDL 1.0 as articulated via three questions raised to successfully deflect an asteroid:

- 1) ‘What is it made of?’
- 2) ‘What shape is it?’
- 3) ‘What is the best choice of technology to make a successful deflection of an asteroid that poses a threat to Earth?’

The FDL has assembled a diverse team of interdisciplinary experts in the field of planetary science, machine learning and deep thinking, drone technology, space-based technology, space mission design and operation to mentor and guide the FDL participants.

Three teams (see Table 1), each composed of 2 planetary astronomers, 2 computer scientists, and a mentor worked together on 6 weeks embedded at the SETI Institute on three applications to answer to the questions above.

FDL Application	Technology	Impact / Ongoing
Teaching a drone to find meteorites in the field / crowdsourcing.	Autonomy / computer vision / Web App	Further improvements in machine vision algorithm
Radically enhancing the process of NEO shape modeling based on radar data.	Unsupervised shape determination based on sparse training data, fast 3d shape generation and Bayesian optimization	White paper and Poster at DPS/EPSC / Ongoing improvements in efficiency. Further development as part of FDL 2.0
Establishing greater confidence on NEO mitigation efficacy	Orbital simulations / trained decision tree	White paper and Poster at DPS/AGU/Vision2050/PDC conference. On going improvement to publish a peer-reviewed article

Table 1: List of applications conducted at FDL 1.0 in 2016, their related technology and the final output of this work

Section II describes the applications conducted by the participants. In Section III we discuss the lessons learned from this first year and Section IV will be an opportunity to discuss the potential of such program in the future.

II. Projects conducted during FDL 1.0

1. Meteorite Recovery.

Fireball tracking camera networks have been established around the world, tracking meteors to determine a source orbit and an impact target region. But to date only 31

meteorites have been found that can be linked to a source orbit. The problem with meteorite recovery is that meteorite searches in target regions can require hundreds of hours to find one fresh fall.

Understanding the composition of asteroids is critical to calculating its potential threat to Earth and to formulating an effective pro-active deflection strategies.

Apart from prohibitively costly sample acquisition missions in outer space, one of the few methods of determining asteroid composition is by studying freshly fallen meteorites and linking them to specific asteroid families. This is achieved by using the fireball trajectories observed as the meteors enter the Earth's atmosphere to compute the orbits of the meteoroids prior to atmospheric entry [2].

During FDL 1.0, a team developed what, to the best of our knowledge, is the first attempt to automate the process of finding small meteorites with unmanned aerial vehicles (UAV) and image processing software.

The team experimented with two types of UAVs: a bespoke radio controlled drone from UVIONIX and a commercially available autonomously controlled drone from 3DR, the Solo. Both drones are quadcopter designs with on board imaging capability. A GoPro HERO 4 was attached to the bottom of the 3DR Solo, capturing an image every second, when flying horizontally at a speed of 3 meters per second. The GoPro was able to fly a grid-search pattern and take still shots at a resolution of 3264×4928 , and we aimed to capture meteorites with minimum resolution of 30×30 pixels, close to a minimum requirement for the image algorithms to detect them.

Freshly fallen meteorites are typically distinguishable from ordinary rocks as they often have a thin (~1-mm) black fusion crust, best recognized when parts of the interior are exposed in places.

The team applied deep learning algorithms to the task of discriminating between images of patches of land with and without meteorites. In particular, they chose to use convolutional neural networks [3], which have achieved state of the art performance in a variety of computer vision based tasks [4]. They trained their algorithm using our data collection of 1-4 cm diameter meteorites photographed in the field, as well as a small number (roughly 280) published photos of meteorites found in relevant terrains. The number of images was artificially increased applying random transformation of the meteorite patch images including rotation, reflection, resolution, brightness and saturation. From this, they built a collection of 32,000 patches of meteorite images (Figure 2).




Figure 2: Examples of data augmentation operations applied to artificially increase our “positive” dataset.

They independently trained 5 convolutional neural network architectures: 3 GoogleNet architectures and 2 AlexNet architectures, using a random subset of 85%

of the labeled data for training and the remainder for validation. The validation accuracies for the 4 models were 99.5%, 99.6%, 99.8%, 99.8% and 99.9%. They created one predictive model by averaging the outputs of each of the 4 neural networks, which achieves a validation accuracy of 99.9%. This method of model averaging helps to remove individual biases that each model may have, improving the final accuracy. To increase the speed of the analysis per patch so it is compatible with a real on-site analysis the team developed alternative models described in [5].



Figure 3: Sample output from the bag of convolutional neural network classification algorithms where the input is an image taken by the UAV. Red boxes are drawn over patches which output 0.75 or higher, and white boxes are drawn over patches with output 0.5 to 0.75. The orange flag in the image marks where the 2 meteorites were placed, both of which were detected in the white boxes. Only 2 boxes are incorrectly labeled out of about 60,000 boxes examined.

Figure 3 is an example of an image taken by the GoPro on the 3DR Solo after being processed by the model. The image contained exactly 2 meteorites which were both found (in the white boxes enlarged to the right of the figure), and rather remarkably, only 2 patches were incorrectly labeled as containing meteorites out of a total of 60,000 boxes examined by the model. This illustrates the potential of convolutional neural network architectures in the task of meteorite detection.

The team believes the primary limitations in achieving even higher accuracy is in the lack of data to train the model and in a suitable pre-selection algorithm to decrease the overall computing time required to examine a field. Over time, as more images of different types of meteorites in various terrains are collected, the algorithms should improve in performance. Realistically, a practical algorithm would require a test accuracy of 99.9%, where an image would have of the order of 10 errors.

With regards to the hardware, there are also improvements that could be made to the UAV. First and foremost, LiDAR software would allow the UAV to maintain a constant height in uneven terrain. Secondly, if the UAV had a fast GPU, it will be possible to identify the meteorites by on-board processing.

2. Accelerated NEO Shape Modeling.

Planetary scientists need precise shape models and extra parameters to compute precise orbits and plan safe robotic landings. Delay-Doppler radar imaging is a powerful technique to obtain information about solar system bodies. Modeling the shape of an asteroid using radar is, however, currently a labor-intensive process, requiring long computer runs and about 4 weeks of human guided iterations to get one asteroid shape.

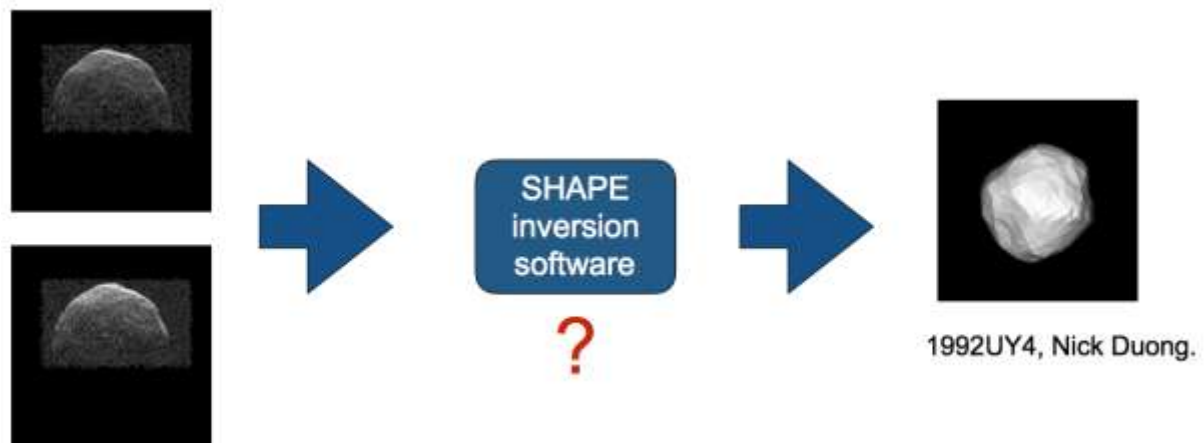


Figure 4: Shape inversion modeling based on radar observation of the asteroid 1992UY4 done in 4 weeks by an undergraduate student at the SETI Institute. Could we improve the speed and automation of this complex data analysis using machine learning algorithm?

Asteroid shapes are critical for asteroid deflection techniques - as any mitigation plan needs to know the three-dimensional form. Should deflection of an impactor not be possible, shape is critical for understanding the potential for damage and planning effective disaster response.

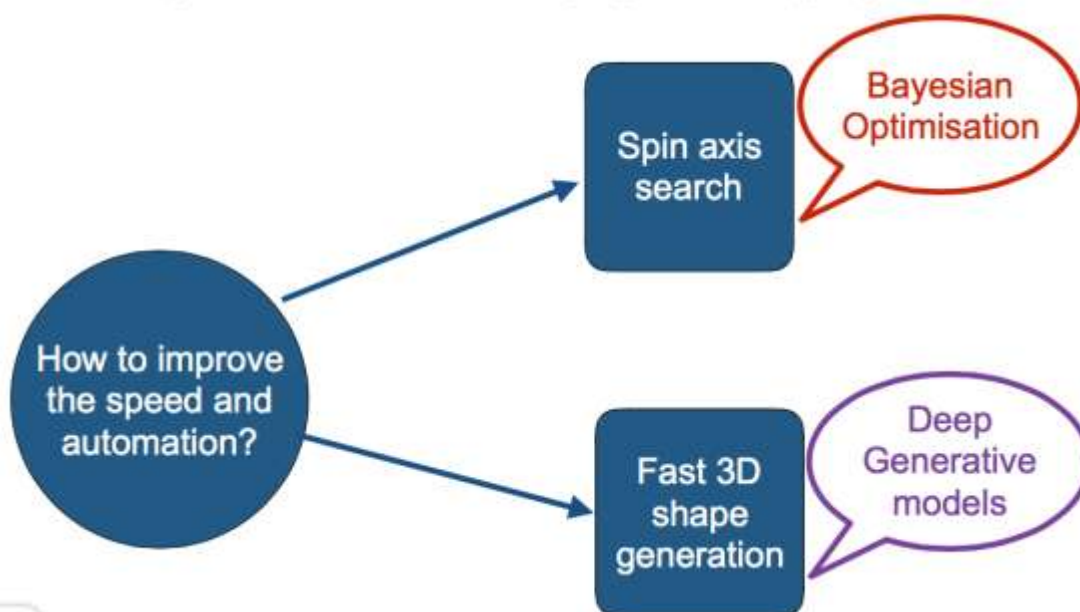


Figure 5: How to improve the speed and automation of asteroid shape modeling in two steps developed by the FDL team. 1) Automated pole direction searches using Bayesian sampling and optimization in conjunction with the SHAPE software 2) Reliable and fast 3D shape generation, separate from SHAPE

A FDL team considered two aspects of delay-Doppler shape inversion with deep generative networks (3D-VAE) and Bayesian optimization.

Bayesian sampling and optimization by automated calls of the existing SHAPE asteroid shape modeling software was tested on radar observations of 1992UY4 and 2000RS11. After 4-6 hours of run time on a single processor, the pole orientations of the asteroids were identified and were the same as those found from human-guided SHAPE fits by under undergraduate students working with the same delay-Doppler datasets over a period of 4 weeks.

To create the 3D shape models of as asteroids the team decided to test deep generative networks for this task. To train the algorithm the team had to create a large database of synthetic data. They used a library of hundreds of shape models from the DAMIT database of asteroid lightcurve shape models database [6] and the JPL database of radar shape models [7]. They then generated 60 sets of simulated Arecibo Observatory radar observations based on the orbits of real near-Earth asteroids as indexed by the JPL Horizons system. Combining 1,620 shape models and those 60 simulated observing sets, they created 546,000 synthetic radar images (including the addition of noise).

The team validated the use of neural network model to derive plausible shape models on this synthetic database. Further develop is required, including developing and testing this algorithm on real radar observations collected at Arecibo and compared with those obtained from current standard shape inversion techniques.

3. NEO Deflector Efficacy

Several technologies have been proposed for impactor deflection, including nuclear explosives, kinetic impactors, and gravity tractors. However, none of these technologies have been developed and fully tested in space. The kinetic impactor method was tested by the Deep Impact mission, which collided a spacecraft with the comet Tempel 1, but the subsequent change in velocity of the comet was not measured. While nuclear explosives are a well-studied technology on the Earth, the testing of nuclear explosives in space was prohibited by the Outer Space Treaty of 1967 [8]. Developing and testing every proposed deflection technology is currently prohibitively expensive. However, if humanity waits until a clear impact threat is detected to select which technologies to use, there may not be time to develop and deploy the chosen deflection technique between the detection of the hazard and the impact. Determining now which technologies are most likely to be useful would allow policy and funding decision-makers to effectively prioritize a subset of the proposed deflection technologies.

Up until FDL, the effectiveness of NEO deflection techniques (such as nuclear, kinetic and gravity tractor) had been studied analytically, but the community lacked a tool that could be used autonomously compare and predict the efficacy of specific techniques.

One of the FDL team has developed a model to map the distribution of parameters of a hypothetical impactor population to the set of technologies that can deflect these objects. This work is described in [9,10] and addresses the following questions:

1. Which deflection method has the highest likelihood of deflecting the broadest range of possible impactors?
2. Which impactor characteristics is the choice of deflection method most sensitive to?
3. Which areas of the impactor parameter space are not covered by current deflection technologies?

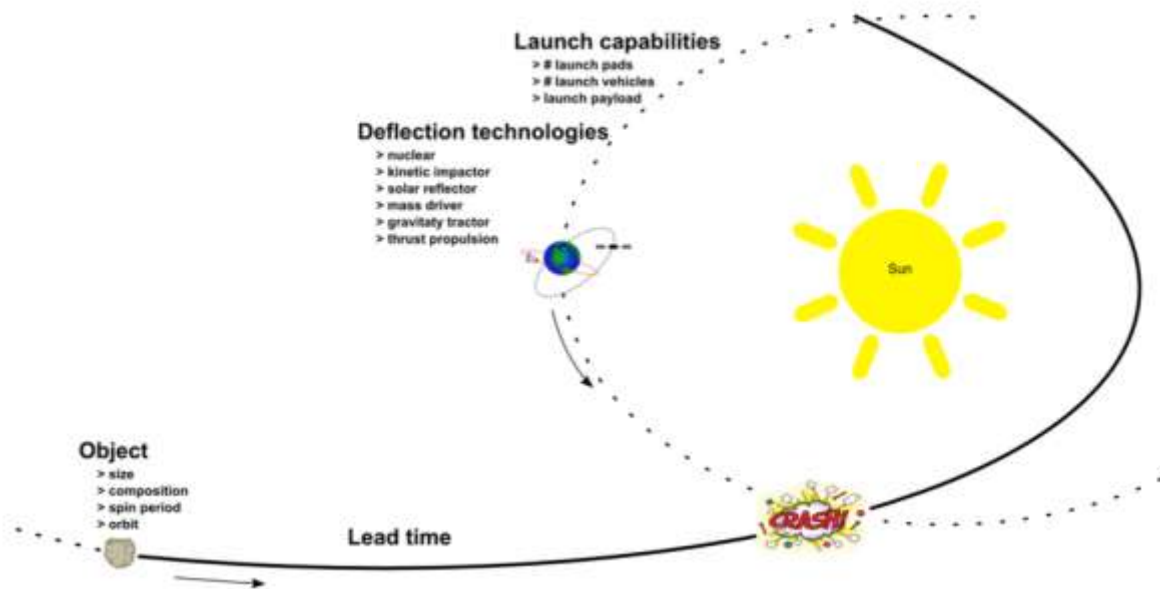


Figure 6: Illustration of the impact scenario and a few of the parameters considered by the model.

The model consists of a machine-learning algorithm that takes as its input the characteristics of a hazardous object (e.g., orbital parameters, size, etc.) and outputs the deflection technologies capable of deflecting the object. To train the algorithm, we produced a set of training data using orbital simulations to simulate the application of a change in velocity, ΔV , to deflect a hazardous object, and a literature search of deflection technologies to calculate which technologies could apply that ΔV , given the object's size. The general considerations and parameters in the orbital simulations and technology calculations are illustrated in Figure 6.

The first step in simulating the efficacy of various deflection techniques is to generate a population of impactors, whose members can then be subsequently deflected. However, the parameters of such a population are not known, since the rate of Earth-impacts is (thankfully) relatively low. Therefore, we generated a simulated population of objects by modifying the orbits of all known Apollo and Aten objects, such that these objects were guaranteed to be Earth-impacting. We performed orbital simulations using the N-body integrator REBOUND, run on the Carnegie Institute of Washington's Memex cluster. Our simulations included the gravitational effects of Jupiter, Venus, and Mars, as well as the Sun and the Earth. We ran ~200 instantaneous-push simulations and ~100 slow-push simulations for each of 8,000

Apollo and Aten orbits that had been altered to impact the Earth. Figure 7 summarizes the results of the instantaneous-push simulations.

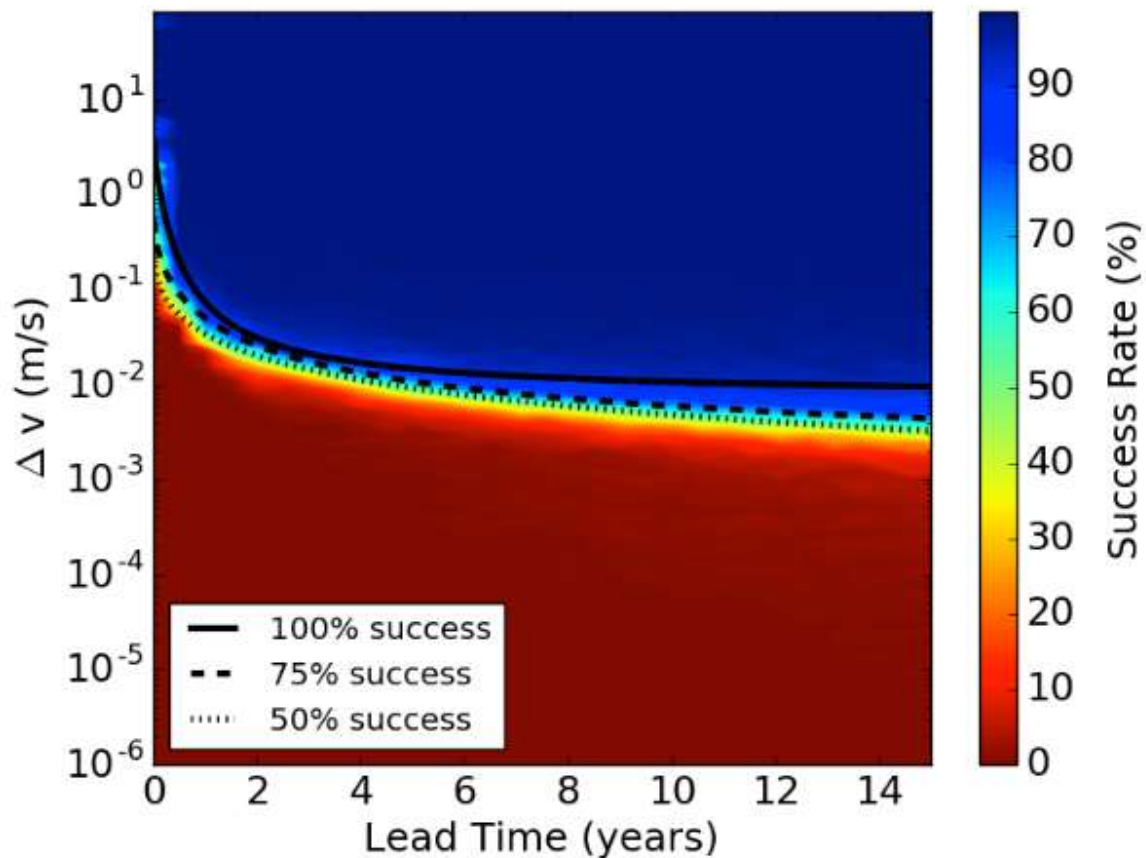


Figure 7: Summary of the instantaneous deflection orbital simulations. The colors indicate the percentage of successful deflections for a given lead time and ΔV applied. Black lines indicate curves fit to the contours for success rates of 50%, 75%, and 100%.

The orbital simulations described above can only reveal which values of ΔV are required to deflect an incoming hazardous object, given its orbit and a lead time. To map these ΔV s to the proposed deflection technologies, we conducted a literature search in order to calculate the ΔV values that each technology can apply, given the object's size. We considered the three most plausible technologies: nuclear explosives, kinetic impactors, and gravity tractors.

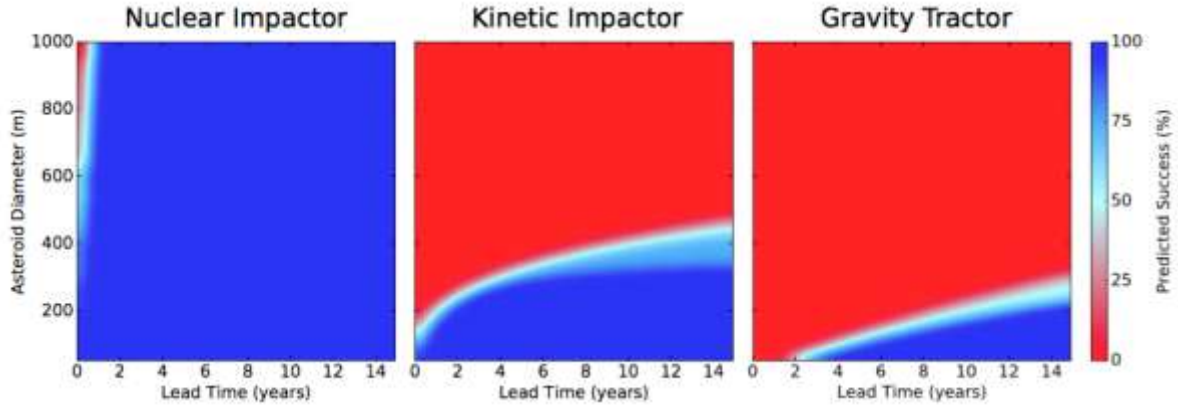


Figure 8: Predicted success rate of each of the three technologies, given an object’s diameter and the lead time between application of the technology and the time of impact.

Figure 8 allows us to directly compare the proposed technologies’ capability of successfully deflecting a hazardous object. From Figure 8 it is clear that the nuclear explosive will be successful in most cases, except for large objects with very short (~ 1 yr) lead times. The kinetic impactor, which imparts generally smaller values of ΔV , is more likely to be successful only for smaller objects (<300 m in diameter) with lead times $> 1 - 2$ yr. The gravity tractor is effective only for the smallest objects with longer lead times. Because the gravity tractor is a slow-push technology, its effects are cumulative over time, and the object diameter for which it is effective increases roughly linearly with time.

Machine learning allows us to produce an algorithm that can predict, given an object’s characteristics, which technologies can successfully deflect the object, much faster than those previously described orbital simulations can be run. The team used the data they produced for the analyses as a training data set for a machine learning algorithm. They chose to train a decision tree algorithm using the scikit-learn package of python, which is effective for classification problems such as this one, and produces a resulting decision tree that can be intuitively understood by the user.

We trained the decision tree on $\sim 80\%$ of the training data, then used the remaining 20% as validation data to test the accuracy of the trained algorithm. We measured the accuracy by inputting the features of the validation data and comparing the decision tree’s output to the labels we had already produced. We performed this cross-validation technique ten times, each time randomly selecting 80% of the data set for training and 20% for validation. The measured accuracy of the trained algorithm was $\sim 98\%$, indicating that this data set is well-suited for classification using the decision tree method. Now that we have trained a decision tree, we can input a set of hazardous objects characterized by orbit, size, β , and lead time, and quickly output the distribution of technologies able to deflect these objects. As observers and theorists continue to refine the simulated distribution of hazardous objects to better reflect reality, we can quickly run the distributions through the decision tree to produce better predictions of which technologies will be most useful. This work has recently been updated by refining the population of impactors and including additional mitigation methods [11].

III. Lessons Learned

1. The success of FDL

As well as make a meaningful contribution to the field of Planetary Defense, FDL itself was a process experiment that explored a number of techniques to accelerate applied research: primarily, the role of interdisciplinary approaches between data science and planetary science researchers, and secondly, how private / public partnerships can tackle both engineering capability gaps and knowledge gaps in the space sciences.

FDL demonstrated promise in a number of dimensions: increasing the likelihood of a useful result and showing how the cadence of beneficial research may be increased by the curation of interdisciplinary teams working in constrained timeframes - while supported by world class experts from both academia and industry.

Moreover it showed that data science / machine learning is a particularly useful catalyst for this kind of format; providing an established and open community of rapidly improving tools and providing the ability to quickly experiment, learn and adapt as needed, quickening a new insight or previously untapped capability.

2. Key learnings from FDL 1.0 (2016) & how to improve the format.

Consultation and exit interviews after FDL 1.0 solicited feedback that was generally very positive. However there were a number of points of learning that can be distilled to the following three insights:

- Where FDL 1.0 could be improved was when activities or advisory placed too much emphasis on analogous or emergent discovery, creating confusion and the perception of sub-optimal use of time. The response to this is to deploy and communicate a more explicit structure for FDL 2.0 in 2017 - in terms of schedule and better targeted external advisory - around "problem space" and "solution space". In other words - invite an expert 'x' who knows about the problem and then expert 'y' who might provide insight on solutions and ensure this is communicated.
- Doctorate level FDL participants are highly competent and there is utility (and cost savings) in getting them to educate each other in each other's domains. Thus in FDL 2.0, time will be scheduled for mutual skills building.
- Whereas FDL 1.0 did a good job of communicating the problem space in application materials, it is recommended that similar material is produced to help planetary scientists understand the potential of machine learning. This is one place where we could leverage the FDL community - in particular the work done in FDL 1.0 and the FDL Fellows.
- The data science / machine learning community is becoming increasingly populated with proprietary processes. Invited contributors with the intention of selling proprietary tools are of little use to the program, which flourishes because of the open nature of many of the most useful platforms and libraries.

IV. Conclusion

FDL demonstrated the potential for an interdisciplinary approach in the field of narrow AI (machine learning and associated techniques such as Deep Neural Nets, Variational Auto Encoders, Generative Adversarial Networks and Bayesian optimization) and planetary defense; reliably surfacing breakthrough applications in accelerated timeframes.

Although capital is required to host the FDL format, it does suggest that breakthroughs and ongoing collaborations useful to the space program can indeed be 'industrialized' using these methods and costs significantly reduced as FDL is scaled and more private sector support is engaged. As point of reference, FDL 2017's six slated projects are each 10% less expensive than 2016 due to economies of scale and increased private sector confidence in the platform. Similar reductions of cost per project are estimated for 2018.

FDL is looking to replicate the success of 2016, by repeating the formula but with more ambitious scope for 2017 [12] (adding teams focused on space weather and space resources) and improvements to the format based on the ideas detailed above. The goal of FDL 2.0 is to further explore applied research useful to the Planetary Defense, Space Weather and Asteroid mining community, through a 'branching' and building on the talent base of the previous FDL - sustaining a community of expertise that is familiar with the interdisciplinary domains, ultimately becoming an association of AI and space science practitioners.

V. References

- [1] NASA Announces Asteroid Grand Challenge.
https://www.nasa.gov/mission_pages/asteroids/initiative/asteroid_grand_challenge.html, 2013.
- [2] P. Jenniskens, A. Rubin and Q. Z. Yin. Fall, recovery and characterization of the Novato L6 chondrite breccia. *Meteoritics & Planetary Science*, 49, 1388–1425, 2014.
- [3] Y. LeCun, L. Bottou, Y. Bengio and P. Haffner. Gradient-Based Learning Applied to Document Recognition *IEEE*, 86(11):2278-2324, 1998.
- [4] J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy and others. Recent Advances in Convolutional Neural Networks. arXiv preprint arXiv:1512.07108, 2015.
- [5] R. I. Citron, A. Shah, S. Siha, C. Watkins, and P. Jenniskens. Meteorite Discovery using an autonomous drone and machine learning. 48th Lunar and Planetary Science Conference, 20-24 March 2017, The Woodlands, Texas. LPI Contribution No. 1964, id.2528, 2017.
- [6] J. Ďurech et al. (2010), DAMIT: a database of asteroid models, *A&A*, 513, A46, available on <http://astro.troja.mff.cuni.cz/projects/asteroids3D/web.php>
- [7] Asteroid Shape Models, available on

<https://echo.jpl.nasa.gov/asteroids/shapes/shapes.html>

[8] United Nations. Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space, including the Moon and Other Celestial Bodies. Adopted and opened for signature, ratification and accession 27 January 1967 , entered into force 10 October 1967 (Article XIV), 1966. 610 UNTS 205.

[9] E. Van Heerden, N. Erasmus, A. H. Greenberg, E. Nesvold, J.-L. Galache, E. Dahlstrom, F. Marchis, The Deflector Selector: A Machine Learning Framework for Prioritizing Deflection Technology Development, American Astronomical Society, DPS meeting #48, id.329.11, 2016

[10] A. H. Greenberg, E. van Heerden, N. Erasmus, F. Marchis, The Deflection Question, AGU Fall Meeting, San Francisco CA, USA, NH13A-1746, December 2017.

[11] E. Nesvold, A. H. Greenberg, E. Van Heerden, N. Erasmus, J.-L. Galache, E. Dahlstrom, F. Marchis, The Deflector Selector: A Machine Learning Framework for Prioritizing Deflection Technology Development, IAA-PDC-17-04-P11, Planetary Defense Conference, Tokyo, May 2017

[12] Frontier Development Lab. [Internet] Moffett Field CA: National Aeronautics and Space Administration Ames Research Center; c2016-17. Available from: <http://frontierdevelopmentlab.org/>