A Self-Supervised Learning-Based Recommender System for NASA's Planetary Data System.

Jeffrey C. Smith*, Douglas A. Caldwell, Matthew S. Tiscareno, Mark R. Showalter, Robert S. French, Mia J. T. Mace, SETI Institute, Mountain View, CA USA, *jsmith@seti.org

Introduction: Many available tools, including the Outer Planets Unified Search [1] (OPUS) hosted by the Ring-Moon Systems (RMS) Node of NASA's Planetary Data System (PDS), allow users to discover data using a variety of metadata parameters. As useful as these tools are, the system requires the user to already have detailed knowledge about what data they wish to obtain. Searching for images, or other data, requires the user to first make hard cuts on the data, then manually sift through many, potentially thousands, of browse products. A method to interactively help direct the user to datasets desired would speed up the search and discovery process.

We set out to develop a proof-of-principle demonstration for a recommender system to be integrated into data discovery tools such as OPUS. We demonstrate the ability to identify a large selection of data desired by the user beginning with just a small set of sample data.

Our vision is to develop a tool that would be useful across PDS, and indeed across NASA's PSD and SMD. We would proceed by continuing to use OPUS as the basis for our proof-of-concept, incorporating into the OPUS website the ability to suggest data products that are similar to a user's current query results, and to do so according to objective ML methods, rather than simply recommending data with similar metadata labels. For example, if a user wants to find craters on outer planet moons, they could provide a very small sample of images of craters and the tool would return a good selection of craters with just a few iterations.

Self-Supervised Learning-Based Recommender: Self-Supervised Learning (SSL) is a machine learning technique where, instead of utilizing a human labeled training dataset to perform supervised learning, the SSL technique probes the data on its own and generates its own self-determined representation of the data. A primary advantage is that the characteristics of the data of interest do not need to be defined *a priori*. It is also possible for the SSL technique to identify characteristics of the data not expected by humans. Our SSL algorithm is based on a method developed by Google Research called SimCLR [2, 3]. This method was adapted by the NASA Frontier Development Laboratory's (FDL) SpaceML team, who developed a tool to take a selection of images and apply the SimCLR method to create an SSL representation [4]. We then adapted and expanded upon the FDL tool for our purposes. Our tool are located at https://github.com/SETI/pds-ml.

Our tool trains an SSL model on a pool of images to be recommended from. The evaluated model output is stored as a hyper-dimensional representation of all the images. A user then provides a small number of sample images. Using a custom semi-supervised learning method, images from the pool are recommended that are similar to those provided by the user. A user provides sample images in two categories: "attractors" and "repulsers." The Recommender then attempts to recommend images similar to the attractors but distant to the repulsers. This technique allows for an iterative approach, where the user begins with a small number of attractors (as few as a single image). The Recommender then returns a user-specified number of recommendations. The user can then either accept the recommendations outright or label the returned images as either good (attractors) or bad (repulsers). The labeled images are then passed back to the Recommender, are combined with the images from the previous iteration and then the Recommender returns a better set of recommended images. This iterative process can continue until the user accepts the recommendations.

After a user supplies their sample images, it is possible to then fully train the SSL model and return a classifier to classify all images in the pool as either similar or not similar to the sample images. This is the technique that is used in the original SimCLR method. After several attempts, we concluded this method does not work well when the variety of pool images is too diverse and the sample images do not represent the full variety of pool images. Instead, we implemented a semi-supervised learning method which utilizes a custom harmonic mean k-nearest neighbors approach.

Example Use Cases: We present here two example use cases. Other tests were performed and, in almost all cases, the method succeeded in recommending the images desired. Performance will certainly vary by the complexity of the problem. Images with very subtle differences will not be as easy to differentiate; however, the example of finding moons within rings demonstrates that subtle features can be identified. There is nothing fundamental to the method that precludes non-image data – any type of data can in principle be used. However, images are easily validated by humans, and thus our examples focus on searching for images in the PDS datasets.

Finding images of Saturn with Specific Ring Orientation. One example use case was to identify images of Saturn with its rings oriented vertically. The goal was to demonstrate not only that a specific object can be identified but also the desired orientation of the object. We began with 8272 Cassini images of Saturn and its moons downloaded from the PDS-RMS node. After generating an SSL representation, we selected 29 images of Saturn with the rings oriented vertically and passed them to the Recommender. On the first iteration, the Recommender returned 100 images with only 11% precision (i.e. 11 of 100 were the desired image and orientation). We then rapidly labeled those 100 images as attractors/repulsers, added them to the original set and passed the new, larger set back to the Recommender. On the second iteration the Recommender returned 100 images with 100% precision (Figure 1). This demonstrates the power of the repulsers in increasing performance.

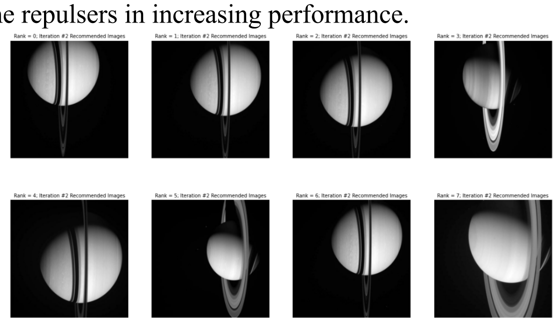

Figure 1: A selection of images of Saturn returned by the Recommender.

Finding Moon-in-Ring Images. A second use case was to identify images with moons embedded in rings from a large set of Cassini images of Saturn's rings. The goal was to demonstrate the ability of the model to identify a subtle feature (tiny moons) in images.

We used 20 images randomly selected from a set of labeled moons (attractors) and no-moons (repulsers) and instructed the Recommender to return recommendations from a set of 707 labeled images. Using the 10 attractors and 10 repulsers resulted in a precision[1] of P=0.83 at a recall[2] of ~0.7. Increasing to 20 attractors and repulsers each gave an average precision of 0.89 and recall of 0.75. See Figure 2 for some example images.

We next ran the Recommender on the full 8,000 SSL training images. All of the top 40 recommendations contained moons, 53 of the top 60 (P = 0.88), and 72 of the top 100 (P=0.72) contained moons. We then looked at 1,985 held-out test set images, in order to validate against our training set, and 15 of the top 15 contained moons, 15 of the top 20 (P=0.75), and 17 of the top 40 (P = 0.43) contained moons.

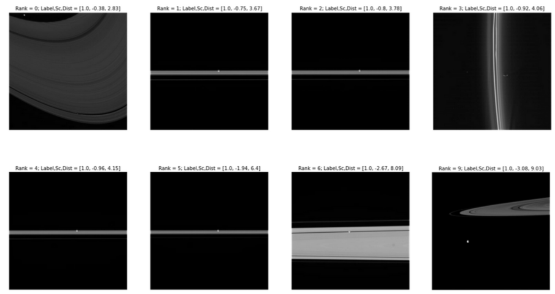

Figure 2: A selection of images of Saturn's rings with embedded moons returned by the Recommender. The moons are the small white dots in the images.

Comparison to Supervised Learning: One of the main advantages of self-supervised learning is it can be used to rapidly train a classifier while requiring only a fraction of the labeling required for supervised learning. To demonstrate this, we compared our Recommender to a dataset developed by another team of a traditional image classification of Cassini images located on the PDS Image Atlas [5]. We downloaded their pool images of Saturn and its moons and computed an SSL representation of the image set. We then used a small random selection of their labeled images as samples for our Recommender. We found that, by using only a very small number of sample images, we could successfully create classifiers with similar performance to their supervised classifiers. For example, by using a mere 10 images classified as "rocks" we could return 150 images with 90% precision to be labeled as "rocks." This is out of 20,994 total images, the vast majority not being of rocks. Likewise, with just 10 samples of "clouds" we could return 800 images with 90% precision to be labeled as "clouds." This demonstrates the power of the SSL recommender method to identify images of interest with similar performance to traditional machine learning methods but using only a small fraction of the human labor required to label images.

Acknowledgments: We thank NASA's SMD Open Source Science Initiative (OSSI) for providing seed money to develop this method.

References: [1] https://opus.pds-rings.seti.org/, [2] Chen T., et al. "A Simple Framework for Contrastive Learning of Visual Representations," International Conference on Machine Learning 2020, https://doi.org/10.48550/arXiv.2002.05709, [3] https://github.com/google-research/simclr, [4] https://github.com/spaceml-org/Self-Supervised-Learner, [5] https://pds-imaging.jpl.nasa.gov/search/

[1] Precision = True Positives / (True Positives + False Positives)
[2] Recall = True Positives / (True Positives + False Negatives)