Onboard Clustering of Aerial Data for Selective Data Return

David S. Hayden
Jet Propulsion Lab, California Institute of Technology
dshayden@asu.edu

Steve Chien
Jet Propulsion Lab, California Institute of Technology
steve.chien@jpl.nasa.gov

David R. Thompson
Jet Propulsion Lab, California Institute of Technology
david.r.thompson@jpl.nasa.gov

Rebecca Castano
Jet Propulsion Lab, California Institute of Technology
rebecca.castano@jpl.nasa.gov

Abstract

Current and future remote space missions, such as the aerial exploration of Titan by an aerobot, have the potential for collecting more data than can be returned for human observation. To relieve this bottleneck, we develop sub-polynomial algorithms for autonomously clustering aerial imagery. Specifically, we construct an Earth-based aerial image dataset as a Titan analogue. We compare performances of k-means-based clusterings against expert clustering. Among 19 low-level image descriptors accounting for color, texture, temporal and spatial arrangement, several are found which allow for clusterings that more closely match that of an expert than do clusterings based on random or periodic sampling. Results show potential for allowing scientists to infer semantic content of all images by downlinking representative images from each cluster. This allows for more efficient use of downlink bandwidth, and therefore higher quality science return in remote space exploration.

1 Introduction

NASA's Solar System Exploration Strategic Roadmap [14] outlines the role of aerial vehicles in the future exploration of the solar system, particularly with respect to missions to Venus or Titan. In the case of Saturn's moon, Titan, an aerobot (blimp) would collect remotely-sensed data some 8km above ground level. It would be capable of circumnavigating the moon within a six month mission. The Cassini-Huygens mission has shown that Titan contains rich and varied landscapes (Figure 1), including smooth and rough terrain, sand dunes, ethane lakes, shorelines, craters, clouds, and possibly cryovolcanoes. With such Earth-like diversity, Titan is of great scientific interest.

Data yield for a Titan mission would be limited not by the rate of image acquisition, but rather by communications constraints. Communication with Earth would be subject to latencies that exceed two hours. Downlink bandwidth is expected to be 4500 bits/second, or 130Mbits / day assuming an 8 hour transmission window [7]. These considerations motivate autonomous methods of classifying aerial image data that could preselect the most scientifically meaningful data for return to Earth.

Previous work in onboard data understanding has focused on rover and satellite platforms. For the Mars Exploration Rovers (MER), science targets such as dust devils or rocks of specific size, albedo, and shape can be
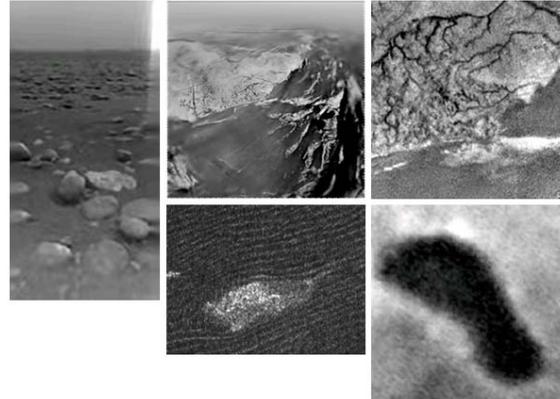

Figure 1: Cassini-Huygens images showing the diversity of Saturn's moon, Titan. Top-row images are from the Huygens probe and show a field containing frozen rocks with horizon, a hill etched by hydrocarbon rain, and part of a dried riverbed, respectively. Bottom-row images are from Cassini radar data and show sand dunes, and hydrocarbon liquid bodies, respectively.

automatically recognized [3]. For the EO-1 Satellite, hazardous events such as fires, floods, and volcanic activities are detected and pertinent data downlinked. [5]. In both cases, targets can be detected without human direction, allowing for automatic data prioritization and improved science return.

Additional challenges exist for an aerobot; it would be in constant motion, but difficult to control due to unpredictable atmospheric currents. Processing would be shared between continuous autonomous control and data processing, but would be limited due to radiation hardening and energy constraints. A typical radiation-hardened processor used in space, the RAD750, is clocked around 200MHz, has 128MB RAM, and can maximally perform 400 MIPS—several orders of magnitude lower than modern desktop computers. These limitations, combined with the diversity of surface features the aerobot might encounter, favors a computationally inexpensive, unsupervised approach that makes few assumptions about the image content the aerobot will encounter. Unsupervised methods have been proposed for selective data return applications in the rover domain. [2, 15]; they have also been widely used for image search and retrieval [4] and image sequence representation [10, 11] Clustering has been applied to aerial imagery [8], though not in an online fashion to our knowledge.

By clustering images as they are collected, an

aerobot mission can analyze greater data volume. For example, the aerobot can downlink the *r* most exemplary images closest to each cluster centroid, providing a broad overview of the types of data collected. If some clusters prove to be consistently interesting then all images from those clusters can be prioritized for downlink. Finally, scientists can also opt to downlink images most dissimilar to the all others—outliers or anomalies. In summary, clustering enables selective data return based on a representative sample, a biased sample, or outliers. These options provide compelling alternatives to other data collection methods such as returning images collected at random, arbitrary or periodic intervals.

In this paper, we investigate the application of unsupervised classification for selective transmission of aerial image data in remote space exploration. We represent images in a metric space to compare their similarities. We identify specific image feature descriptors to encourage clusters based on semantic content such as presence of horizons, clouds, and water bodies. A broad survey of different image features suggests several that are both computationally efficient for spacecraft computing resources and relevant to the image categories identified by planetary scientists.

Section 2 begins by describing a small Earth-based aerial image dataset. Section 3 details the low-level, computationally inexpensive features we consider. These features account for color, texture, spatial arrangement, and time. Section 4 details both our algorithmic and manual (expert) clustering methods, as well as the metric with which we compare them. Sections 5 and 6 compare the effectiveness of different feature and clustering parameters. For our dataset, clustering based on several frequency-space features can more closely match an expert's clustering than periodic or random sampling.

2   Dataset

We constructed a dataset of aerial imagery using a consumer-grade digital camera (Canon PowerShot SD850 IS) with resolution limited to 1600 x 1200. A total of 162 images were collected during a commercial flight from New York to Los Angeles. Images primarily contain shots dominated by sky, horizon, or undeveloped land. Some contain clouds, discernable water bodies, developed land, or have small portions of the plane's wing or window at one more edges. Example images are displayed in section 4.

Factors such as varied terrain, presence of clouds, horizons, water bodies, and artifacts (e.g. the occasional window or wing obstruction) made this dataset a particularly appropriate Titan analogue.

3   Features

We chose features to represent the color, texture, time, and spatial arrangement of each image. In order to reduce computational costs we favored simple features based on first- or second-order statistics that required little preprocessing.

We split the features into four themes: edge, color, frequency, and time. The edge and frequency features correlate with image texture, color captures basic color statistics, and time is an integer denoting the temporal order of image acquisition.

3.1 Edge Features

Let I be an $m \, x \, d \, x \, 3$ image, and I' be its $m \, x \, d$ grayscale. Let I'' be the resulting binary image from performing convolution with a Sobel operator on image I'. Let G_y, G_x be the $m \, x \, d$ matrices representing vertical and horizontal gradient responses, respectively. Let $\nabla L = \sqrt{G_y^2 + G_x^2}$ be the $m \, x \, d$ matrix representing gradient magnitudes, and $\theta = atan2(G_y, G_x)$ be the $m \, x \, d$ matrix representing gradient orientations. Then the edge features are:

$$F_0 = \frac{1}{md} \sum I' \quad \text{(edge density)}$$

$$F_1 = \frac{1}{md} \sum (\nabla L) \quad \text{(mean gradient magnitude)}$$

$$F_2 = H(\nabla L) \quad \text{(gradient magnitude entropy)}$$

$$F_3 = H(\theta) \quad \text{(gradient orientation entropy)}$$

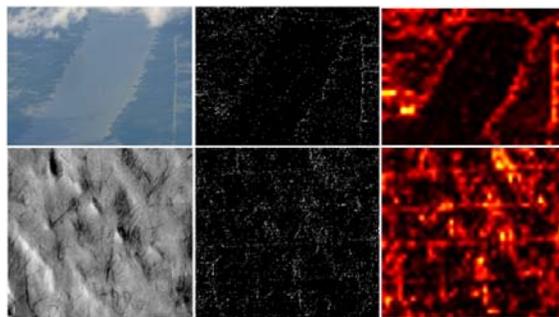Figure 2 contrasts the edge densities of images with smooth and rough terrain.

Figure 2: The density of edges (middle) as visualized in a heat map (right) provide texture information. Water bodies (top) tend to be homogenous regions containing few edges as compared to purely land-based imagery, such as this dust devil-torn landscape (bottom) in the southern hemisphere of Mars taken by THEMIS (image ID V07829003, [18]). The edge density of the top image is 0.02 while the edge density for the bottom is 50% greater at 0.03.

3.2 Color Features

Let P_i be the $m \times d$ matrix of pixels in band i of I. Then, for all P_i the color features are:

$$F_{4..6} = \text{standard-deviation}(P_i)$$

$$F_{7..9} = \text{mean}(P_i)$$

$$F_{10..12} = \min(P_i)$$

$$F_{13..15} = \max(P_i)$$

Figure 3 contrasts the color features of two images containing different terrain.
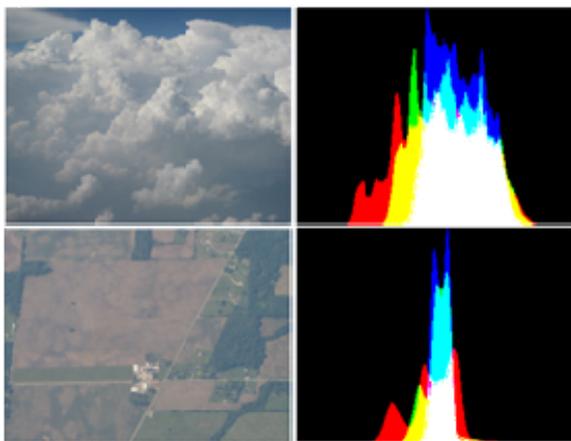

Figure 3: Terrestrial images (left) and their color histograms (right). Images dominated by clouds or sky will have histograms distinctive from many ground-based images (bottom). The top image has greater expression of all colors (standard deviation about the mean of red, green, and blue bands are 38, 30, 23) resulting in an image dominated by white. The bottom image contains diminished color expression as demonstrated by the tighter standard deviation of colors about their means (22, 12, 8).

3.3 Frequency Features

Let f be the resulting $m \times d$ matrix after the 2D Fourier transform on I', and let $G(f) = |f|^2$ be the $m \times d$ power spectrum of f. For the normalized power spectrum $N = \frac{G(f)}{\sum G(f)}$ the frequency features are given by:

$$F_{16} = \sum_{x>0, y>0} N \qquad \text{(quadrant-1 energy)}$$

$$F_{17} = \sum_{x<0, y>0} N \qquad \text{(quadrant-2 energy)}$$

These features were motivated by [12], which were found by a survey of 28 low-level frequency-space statistics to best discriminate a subset of the Brodatz textures.

Figure 4 provides visualization of a 2D Fourier transform in terms of its amplitude on both real and synthetic data.
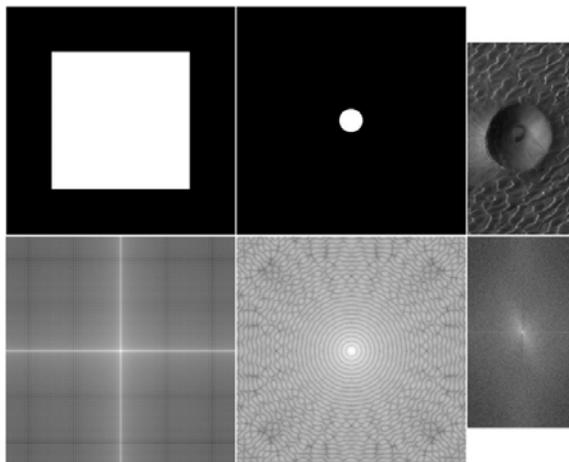

Figure 4: Top row shows images of a square, circle, and sand dunes within Mars Proctor Crater (image ID V02571003 [18]). Bottom row shows respective log-scale magnitude images of their Fourier transforms. Notice that quadrants 1 and 3, as well as quadrants 2 and 4, of the magnitude images are symmetric.

3.4 Time Features

For a sequence of images collected in serial, let the acquisition order I_t be an integer representing the ordering of each image:

$$F_{18} = I_t \qquad \text{(acquisition order)}$$

3.5 Spatial Features

In addition to accounting for texture, color, and time, we attempt to capture the spatial distribution of image features. Traditional approaches use image segmentation or region-growing methods to find objects or areas of uniform color or texture. These methods may be too expensive for aerobot applications.

Instead, we account for spatial arrangement by splitting each image into $n \times n$ equal-sized subimages and collecting the same features in each. Hence, for $n = 1$, we collect each desired feature once on the whole image, whereas for $n = 4$ we collect each feature in each of 16 subimages (except for acquisition order, which is never collected more than once). These subimage features are appended into an ordered vector whose dimensionality grows in proportion with the number of subimages.

4   Clustering

We cluster using iterative k-means per the standard Lloyd's algorithm [13]. We initialize cluster centers using random datapoints and use a Euclidean distance metric to assign cluster membership. Image features were standardized prior to clustering.

Distances in the Euclidean space are sensitive to noisy or redundant dimensions. To account for these factors we employed a linear dimensionality reduction using Principal Component Analysis (PCA) with varying numbers of principal components.

4.1 Expert Labeling

We elicited a manual clustering of our dataset from a planetary volcanologist to serve as a ground-truth standard for evaluation. The expert's only formal introduction to the task was a three minute introduction to custom software that facilitated manual clustering through a graphical, drag-n-drop interface. Additionally, the following written prompt was provided:

> Suppose that the following aerial images were taken of an environment for which we have little knowledge or data. Furthermore, suppose that you may not be able to receive all images. Please sort these images into 5 groups in such a way that if you could only receive a small number of images from each group, you could reasonably infer the content of the remaining images in that group.

We considered allowing the expert(s) to choose the number of clusters, but decided that this would provide too much flexibility given our intentionally vague prompt. This could create confusion about clusters' appropriate extent and roles. It is especially important since accepted methods of interpolating multiple, non-homogenous expert clusterings (e.g. [16]) are inadequate if scientists pursue multiple, independent goals. We settled on 5 clusters for the simple reason that the number of clusters should be at least an order of magnitude less than the size of our dataset. The current study uses the data provided by a single expert. Upon interview, the expert felt that five was a mostly adequate number; six would have been ideal so that an outlier group could have been established. Notably, the expert spent roughly 20 minutes clustering 162 images.

Figure 5: Five randomly chosen images from each of five clusters created and named by a planetary volcanologist

The expert sorted images by the semantic distinctions "Rivers", "More Clouds", "Land", "Horizon", and "Desert." Figure 5 shows representative images selected at random from each. In terms of low-level features, images in the "Horizon" category contain a line separating the image into two regions of color and texture: above the horizon the sky contains nearly uniform texture and a gentle gradient from light to dark blue, while below the line both texture and color dramatically vary. "Land" images largely contain patches of ground in shades of brown and green while "Desert" images are dominated by only brown patches. Texture in all images generally becomes smoother as the altitude increases. The clouds in "More Clouds" images contain patches different from neighboring areas in color and texture. "Rivers" images appear to be the most difficult to describe in image primitives. One observation is that the rivers are meandering polylines of widely varying thickness and are typically of a different color and smoother texture than their surroundings.

4.2 Cluster Comparison

We compare automatic clusterings against the expert standard using the information theoretic adjusted mutual information (AMI), which we briefly derive here.

Given dataset $S = \{s_1, s_2, \ldots, s_N\}$ and clusterings

$$U = \{U_1, U_2, \ldots, U_R\} \qquad (1)$$

$$V = \{V_1, V_2, \ldots, V_C\} \qquad (2)$$

where $\bigcap_{i=1}^{R} U_i = \emptyset$, and $\bigcup_{i=1}^{R} U_i = S$ (e.g. the U_i are a partitioning, or clustering, of S, and similarly for V). Then, the probability that a random data $s \in S$ is also contained in some cluster U_i is

$$P_u(i) = \frac{|U_i|}{N} \qquad (3)$$

The probability that s is contained in some V_j is

$$P_v(j) = \frac{|V_j|}{N} \qquad (4)$$

The joint probability that $s \in U_i$ and $s \in V_j$ is

$$P(i,j) = \frac{|U_i \cap V_j|}{N} \qquad (5)$$

The mutual information is then defined as

$$MI(U,V) = \sum_{i=1}^{R}\sum_{j=1}^{C} \qquad (6)$$
$$P(i,j)\log\left(\frac{P(i,j)}{P_u(i)P_v(j)}\right)$$

Mutual information quantifies how much knowing about one clustering tells us about the other. Though it is symmetric and non-negative, it is not upper-bounded by a constant, and so is not useful as a general metric for comparing clusterings. Furthermore, Vinh et al. demonstrate that mutual information does not take a constant value when comparing random clusterings, and tends to grow with the number of clusters [17]. They use a hypergeometric model of randomness to derive an expected value for two random clusterings. This permits a correction similar to the Adjusted Rand Index [9] that ensures random clusterings produce a constant value, This correction yields the Adjusted Mutual Information (AMI):

$$AMI(U,V) = \frac{(MI(U,V) - E\{MI(U,V)\})}{\max\{H(U),H(V)\} - E\{MI(U,V)\}} \qquad (7)$$

The entropies of clusterings U,V denote the uncertainty in a data point's cluster membership:

$$H(U) = -\sum_{i=1}^{R} P(i)\log P(i) \qquad (8)$$

$$H(V) = -\sum_{j=1}^{C} P(j)\log P(j) \qquad (9)$$

The denominator in AMI corrects for randomness and serves as a normalization, as otherwise $MI(U,V) \le \min(H(U),H(V))$. Furthermore, $AMI(U,V) = 0$ only when equal to its expected value (e.g., that expected by comparing two random clusterings), and $AMI(U,V) = 1$ when clusterings U,V are identical.

5   Results

5.1 Exhaustive Parameter Search

We begin with an exhaustive search over the parameter space to identify promising features. Specifically we cluster the dataset once for each possible non-empty subset of the four feature themes (edge, color, frequency, time). We consider each theme as the smallest unit (rather than each individual feature), which yields 15

potential combinations. Finally, we use subimage decompositions of size $1, 4, 9,$ and 16, corresponding to $n = 1, 2, 3,$ and 4.

We apply PCA with different dimensionalities to account for potential redundancy within, or interplay between, feature themes (e.g. the possibility that one edge feature and one color feature are most discriminative). We vary the number of principal components used from $1, 2, ..., 19, 20, 24, 28, ..., 96, 100$ or up to the dimensionality of the base feature set. Finally, we also vary the number of clusters, k, over $3, 5, 7, 9, 11$.

For each set of parameters, we compare algorithmic clustering with the expert clustering. In figure 6, the maximal AMI for each feature theme across all parameters is reported. For convenience, we summarize the parameters:

- The 15 non-empty subsets of themes: {edge, color, frequency, time}.
- Using PCA vs. not using PCA.
- When using PCA, the number of principal components from $1, ..., 100$.
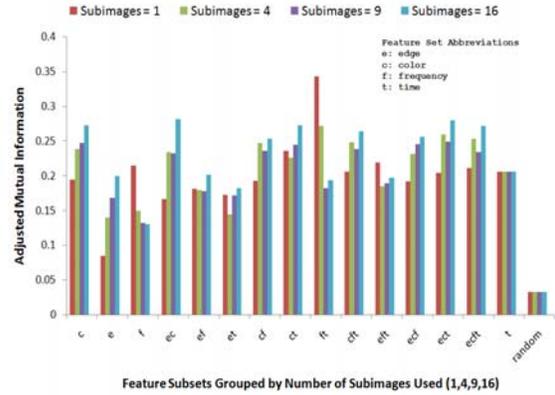- Number of clusters, k, from $\{3,5,7,9,11\}$.

Figure 6: Best clustering performances among feature themes over increasing numbers of subimages. Reference performances of random and acquisition order (t) clustering are also displayed to the far right.

Increasing the number of subimages improves the AMI for all theme combinations that do not make use of frequency-space features. Except for the FT features, AMI does not exceed 0.26 for any subset of features unless collected under 16 subimages, at which point dimensionality reduction becomes necessary. For one subimage, the FT features yield the highest overall AMI of 0.343. We exclusively consider this combination in the analysis that follows because it is fast to collect and data are represented in a mere three dimensions.

5.2 Frequency-Space and Time Features

In further investigating clusterings with FT features,

we first verify the trend in figure 6 that larger numbers of subimages will lower performance. We fix the number of clusters to $k = 5$ and conduct 100 trials for each of 1, 4, 9, and 16 subimages, and over all possible number of principal components. For each variation in the number of subimages we restrict our results to the set of trials that yielded the highest AMI over all possible number of principal components. Figure 7 displays these results.
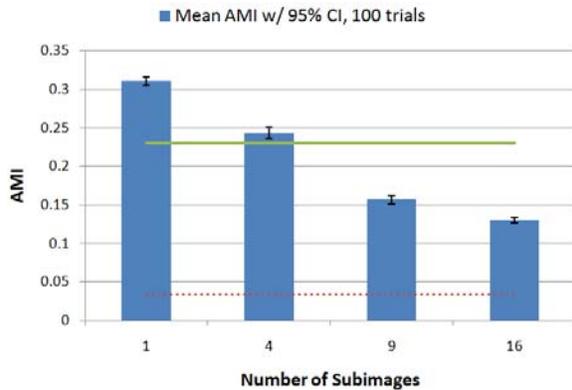

Figure 7: FT-based clustering performance with varying subimages. The solid and dotted horizontal lines represent mean AMI over 100 trials of acquisition order and random clustering, respectively.

We next consider the utility of PCA for one subimage. This is an appropriate question since the FT combination may not require all features for optimal performance.
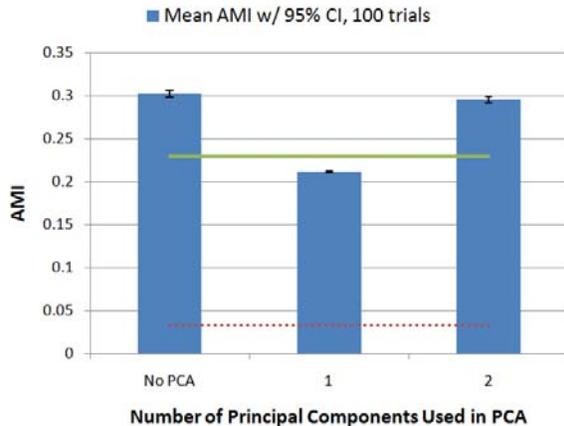

Figure 8: FT-based clustering performance using 5 clusters and 1 subimage while varying the number of principal components. The solid and dotted horizontal lines represent mean AMI over 100 trials of acquisition order and random clustering, respectively.

We see that there is no statistical significance between the base feature set and a compressed feature set

(Figure 8). Considering the performances of acquisition order clustering and frequency-space clustering in figure 6, it may be that only one frequency-space feature, paired with acquisition order, is necessary. Alternatively, designers can implement all three features, opt out of dimensionality reduction, and suffer no detriment.

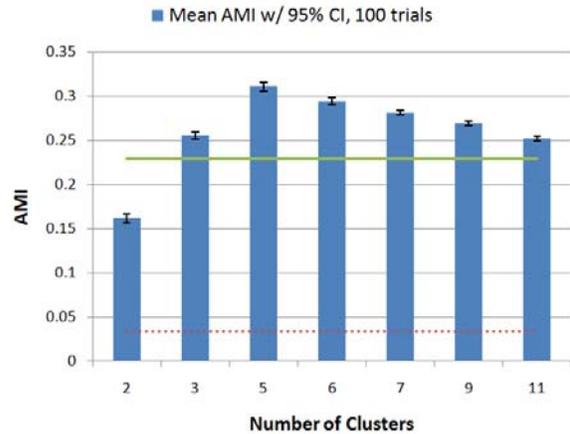Next, we vary the numbers of clusters within 1 subimage.

Figure 9: FT-based clustering performance using 1 subimage and no PCA while varying the number of clusters, k. The solid and dotted horizontal lines represent mean AMI over 100 trials of acquisition order and random clustering, respectively.

Figure 9 shows optimal performance when the number of clusters matches the number used in the expert's manual evaluation.

Finally, we compare the results of clustering with FT features against methods based on acquisition order and random clustering in figure 10.
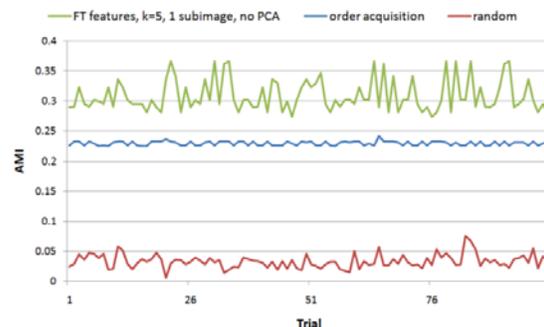

Figure 10: Comparison of FT-based clustering performance over 100 trials against acquisition order and random clusterings

Random clustering results in a low, fixed performance score and clustering with frequency and

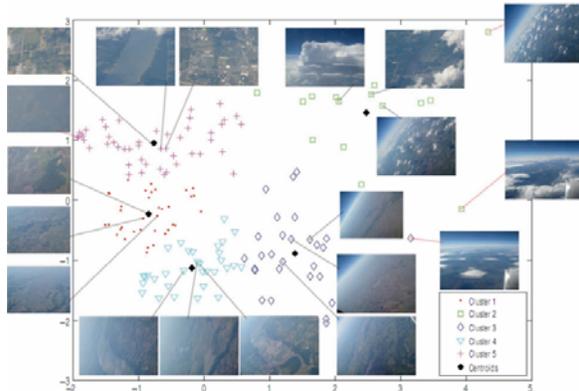time features consistently outperforms clustering based on time alone.

Figure 11: Example algorithmic clustering of a terrestrial dataset using FT features. Three representative images closest to each cluster centroid are displayed (gray lines). Four outliers are also displayed (red lines).

More qualitatively, Figure 11 displays an FT clustering in two dimensions with near centroid and outlier images. Most striking is the separation between images containing horizon and/or clouds from those containing ground-based images. Cluster 2 favors images with clouds while cluster 3 favors images with horizon. Both horizon and clouds are abundant in the outlier images contained by Cluster 2. Both of these clusters contain images with significantly lower energy in the power spectrum of their Fourier transforms when compared to the rest of the images. This is likely due to the contribution of strong, lower-frequency signals from homogenously textured clouds and sky, and the weakly periodic nature of images speckled with clouds. There is no clear distinction between Clusters 1, 4, and 5 in terms of the expert categories, but it is interesting to note that the representative images in Cluster 5 have highly detailed textures and are taken at a relatively low altitude.

6 Discussion

Our initial investigation revealed three features which outperform both random and acquisition order-based clustering, which respectively serve as analogues for random sampling and periodic sampling in an environment. These features are relatively fast and simple to collect. The requisite Fourier transform is an $n \log n$ operation [6] while the normalization and energy sums are linear with respect to image size. Dimensionality reduction need not be used. With a fixed number of iterations, k-means runs in linear time. Therefore the total algorithmic complexity is sub-polynomial in image size and linear in the number of images.

The computational cost of the Fourier transform may be reduced if, in future work, it were found that these features worked as effectively on thumbnail-sized images or independent subimages.

It is well known that solutions provided by k-means are sensitive to initialization. In our results, we see that performance can vary by as much as 0.05 AMI, but it never underperforms clustering based on acquisition order. Even so, there has been significant work on clever initialization of k-means that tends to improve both performance and runtime (e.g. [1]). These should be considered for future work.

Exploration of domain-specific features would also be useful. Determining methods that reliably form clusters for specific terrain (or its lower-level counterpart, texture) could aid specific science objectives. Beyond forming clusters based on terrain, it might also be interesting to consider forming clusters based on raw image quality. That is, can we easily set aside images which are marred with noise or sensor artifacts, or were collected with unfocused optics?

Finally and most importantly, we are comparing against a single expert and dataset. These results should be validated against additional experts and datasets.

7 Conclusion

We have explored onboard, computationally inexpensive clustering for improving the science return of missions where more data can be collected than returned for human observation. Motivated by the potential for the aerial exploration of Titan via an aerobot, we have collected an Earth-based aerial image dataset and compared k-means clustering with that of a planetary volcanologist. Among 19 low-level features accounting for color, texture, spatial and temporal arrangement, three were found which more closely match the manual clustering of an expert than do clusterings formed by random or periodic sampling.

8 References

[1] D. Arthur and S. Vassilvitskii, "k-means++: the advantages of careful seeding," SODA '07: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms, New Orleans, Louisiana: Society for Industrial and Applied Mathematics, 2007, pp. 1035, 1027.

[2] R. Castano, T. Estlin, R.C. Anderson, D.M. Gaines, A. Castano, B. Bornstein, C. Chouinard, and M. Judd, "Oasis: Onboard autonomous science investigation system for opportunistic rover science: Research Articles," J. Field Robot., vol. 24, 2007, pp. 379-397.

[3] R. Castano, T. Estlin, D. Gaines, A. Castano, C. Chouinard, B. Bornstein, R. Anderson, S. Chien, A. Fukunaga, and M. Judd, "Opportunistic rover science: finding and reacting to rocks, clouds and dust devils," Aerospace Conference, 2006 IEEE, 2006, p. 16 pp.

[4] Y. Chen, J.Z. Wang, and R. Krovetz, "CLUE: Cluster-based Retrieval of Images by Unsupervised Learning," IEEE Transactions on Image Processing, vol. 14, 2003, p. 2005.

[5] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, D. Boyer, "Using Autonomy Flight Software to Improve Science Return on Earth Observing One," Journal of Aerospace Computing, Information, and Communication , April 2005.

[6] J.W. Cooley and J.W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series," Mathematics of Computation, vol. 19, Apr. 1965, pp. 297-301.

[7] J. Hall, V. Kerzhanovich, A. Yavrouian, J. Jones, C. White, B. Dudik, G. Plett, J. Mennella, and A. Elfes, "An aerobot for global in situ exploration of Titan," Advances in Space Research, vol. 37, 2006, pp. 2108-2119.

[8] K.R. Harvey and G.J.E. Hill, "Vegetation mapping of a tropical freshwater swamp in the Northern Territory, Australia: a comparison of aerial photography, Landsat TM and SPOT satellite imagery," International Journal of Remote Sensing, vol. 22, 2001, pp. 2911–2925.

[9] L. Hubert and P. Arabie, "Comparing partitions," Journal of Classification, vol. 2, Dec. 1985, pp. 193-218.

[10] L.S. Kennedy and M. Naaman, "Generating diverse and representative image search results for landmarks," Proceeding of the 17th international conference on World Wide Web, Beijing, China: ACM, 2008, pp. 297-306.

[11] S. Krishnamachari and M. Abdel-Mottaleb, "Image browsing using hierarchical clustering," Proceeding of the IEEE Symposium on Computers and Communications (ISCC), 1999.

[12] S. Liu and M.E. Jernigan, "Texture analysis and discrimination in additive noise," Comput. Vision Graph. Image Process., vol. 49, 1990, pp. 52-67.

[13] S. Lloyd, "Least squares quantization in PCM," IEEE Transactions on Information Theory, vol. 28, 1982, pp. 129–137.

[14] NASA. "2006 solar system exploration roadmap" Technical report, NASA, 2006. Washington, DC.

[15] D. Thompson, T. Smith, and D. Wettergreen, "Information-optimal selective data return for autonomous rover traverse science and survey," Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on, 2008, pp. 968-973.

[16] A. Vailaya, A. Jain, and Hong Jiang Zhang, "On image classification: city vs. landscape," Content-Based Access of Image and Video Libraries, 1998. Proceedings. IEEE Workshop on, 1998, pp. 3-8.

[17] N.X. Vinh, J. Epps, and J. Bailey, "Information theoretic measures for clusterings comparison: is a correction for chance necessary?," Proceedings of the 26th Annual International Conference on Machine Learning, Montreal, Quebec, Canada: ACM, 2009, pp. 1073-1080.

[18] Christensen, P.R., N.S. Gorelick, G.L. Mehall, and K.C. Murray, THEMIS Public Data Releases, Planetary Data System node, Arizona State University, <http://themis-data.asu.edu>.