

A Contrastive Learning Approach to Auroral Identification and Classification

1st Jeremiah W. Johnson

*Dept. of Applied Engineering & Sciences
University of New Hampshire
Manchester, NH, USA
jeremiah.johnson@unh.edu*

2nd Swathi Hari

*Dept. of Applied Engineering & Sciences
University of New Hampshire
Manchester, NH, USA
spj322@wildcats.unh.edu*

3rd Donald Hampton

*Geophysical Institute
University of Alaska–Fairbanks
Fairbanks, AK, USA
dhampton@alaska.edu*

4th Hyunju K. Connor

*Geophysical Institute
University of Alaska–Fairbanks
Fairbanks, AK, USA
hkconnor@alaska.edu*

5th Amy Keesee

*Department of Physics
University of New Hampshire
Durham, NH, USA
amy.keesee@unh.edu*

Abstract—Unsupervised learning algorithms are beginning to achieve accuracies comparable to their supervised counterparts on benchmark computer vision tasks, but their utility for practical applications has not yet been demonstrated. In this work, we present a novel application of unsupervised learning to the task of auroral image classification. Specifically, we modify and adapt the Simple framework for Contrastive Learning of Representations (SimCLR) algorithm to learn representations of auroral images in a recently released auroral image dataset constructed using image data from Time History of Events and Macroscale Interactions during Substorms (THEMIS) all-sky imagers. We demonstrate that (a) simple linear classifiers fit to the learned representations of the images achieve state-of-the-art classification performance, improving the classification accuracy by almost 10 percentage points over the current benchmark; and (b) the learned representations naturally cluster into more clusters than exist manually assigned categories, suggesting that existing categorizations are overly coarse and may obscure important connections between auroral types, near-earth solar wind conditions, and geomagnetic disturbances at the earth’s surface. Moreover, our model is much lighter than the previous benchmark on this dataset, requiring in the area of fewer than 25% of the number of parameters. Our approach exceeds an established threshold for operational purposes, demonstrating readiness for deployment and utilization.

Index Terms—Unsupervised learning, representation learning, heliophysics, aurora.

I. INTRODUCTION

Understanding the interaction of the solar wind with the Earth’s magnetosphere and upper atmosphere is one of the key topics in heliophysics [1], [2]. Aurora Borealis and Aurora Australis are formed when high-energy electrons and ions precipitate from the magnetosphere to the upper atmosphere due to the coupling of solar wind and magnetosphere. These auroral features at a scale of 10 – 10,000km hold the key information of the magnetospheric coupling processes at a scale of 1 – 1,000 earth radii. Since it is difficult to cover the entire

This work is supported by the National Science Foundation grant OIA-1920965.

magnetosphere with a limited number of satellites, ground-based auroral images have been critical to understanding the propagation of the solar wind energy to near-Earth space [3]–[5].

With decades of ground and space observations, global aurora patterns of 1,000 – 10,000km and their general mechanisms are well understood [6]. However, the morphology of smaller-scale aurora forms (less than 1000km) and their connections to mid- to large-scale magnetospheric dynamics (less than 100 earth radii) are still under question, due both to the complexity of aurora features and the abundance of aurora images collected. A better understanding of local-scale auroral morphology is critical to improving our understanding of the solar wind–magnetosphere–upper atmosphere interaction.

The heliophysics community has over decades amassed vast collections of images of Aurora Borealis and Aurora Australis. To date, the vast majority of this data is unclassified. Classifying this data is a first and important key step to identify local-scale auroral features and thus to enable a deeper analysis on the link between auroral features and magnetospheric dynamics.

II. RELATED WORK

Automatic classification of auroral images has been studied in the past; tradition computer vision approaches reliant on the construction of hand-designed features include [7]–[10]. While most of these approaches are only effective when limited to binary *aurora/no aurora* classification, the approach of [11] using a hidden Markov model and incorporating temporal dynamics is able to identify four distinct categories of aurora with a positive detection rate of up to 85%.

Computer vision is now dominated by convolutional neural network (CNN)-based algorithms that have achieved remarkable performance on a wide range of challenging tasks. Following the general trend in computer vision, there have

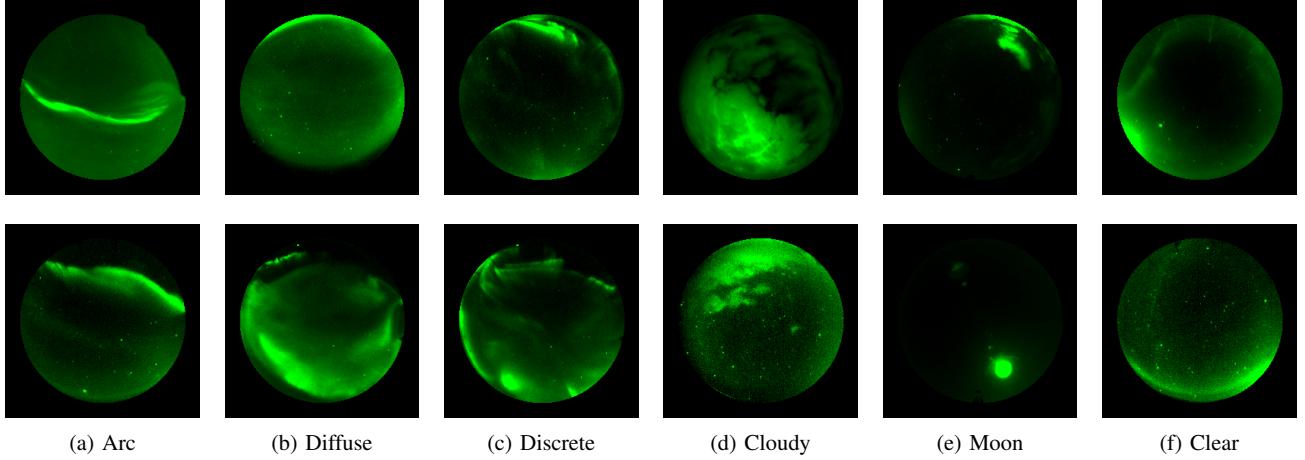


Fig. 1: Random samples of images from the OATH dataset showing the distinct categories of aurora.

been several recent efforts to use CNN-based models for automatic auroral classification. In [12], Clausen *et. al.* manually label 5,824 images of aurora obtained from THEMIS all-sky imagers into the distinct categories *Arc*, *Discrete*, *Diffuse*, *Clear/No Aurora*, *Cloudy*, and *Moon*. Using an Inception model [13] pretrained on the ImageNet dataset [14], [15], they extract feature vectors and fit a ridge classifier to them to obtain an 82% accuracy. In [16], the authors manually classified 3,846 auroral images into the highly imbalanced categories *Arc aurora*, *Auroral breakup*, *Colored aurora*, *Discrete aurora*, *Edge aurora*, *Faint Aurora*, *Patchy Aurora*, then compare classification results from a range of machine learning algorithms, including CNN-based architectures.

There are two closely related challenges inherent to using conventional CNN-based approaches to auroral image analysis. First, supervised learning methods require a substantial quantity of unambiguous, high-quality labeled data be available for training. While there exists massive amounts of high-quality auroral image data, to date very little of this data has been labeled. The use of ImageNet-pretrained CNNs as in [12] presents a possible alternative, as pretraining is often effective at reducing the amount of data necessary to train a CNN to convergence [17], but pretraining only partially addresses the challenge of limited data and may introduce other issues and biases [18]. Secondly and more importantly, manually labeling auroral images itself presents additional challenges beyond the time and expense required: labels for auroral images are often chosen based on subjective, qualitative judgements, and there is frequent disagreement between researchers regarding what an appropriate label set is. Complicating the picture further is the fact that even after a label set is decided on it is common for all-sky auroral images to contain more than one type of aurora, raising the additional question of how to decide what the correct classification should be.

In the past few years, unsupervised machine learning algorithms such as [19], [20], [21] have begun to achieve levels of performance rivaling those of their supervised counterparts

on benchmark tasks and datasets, including CIFAR10, CIFAR100, and ImageNet. However, the utility and efficacy of recent unsupervised models for practical tasks and real-world data is not yet well established. Unsupervised models afford an opportunity to address the challenges just described in auroral imaging simultaneously: the learned representations of auroral images may be useful for quantitatively determine meaningful labels, and can be used to apply such labels automatically, facilitating the creation of much larger labeled auroral image datasets. Our contribution in this work is as follows:

- We modify and adapt the *Simple Framework for Contrastive Learning of Representations* (SIMCLR) model introduced in [19] and apply it to a publicly available auroral image dataset released with the publication of [12], consisting of 5,824 images of aurora obtained from THEMIS all-sky imagers [22] and manually categorized into the distinct categories *Arc*, *Discrete*, *Diffuse*, *Clear/No Aurora*, *Cloudy*, and *Moon*. Our approach leads to sufficiently informative representations to enable a simple linear classifier to obtain state-of-the-art classification performance across these categories, reducing the top-1% error rate from the previous benchmark by almost 10 percentage points. Moreover, our model obtains this performance while requiring less than 25% of the number of parameters of the model used to obtain the previous state-of-the-art.
- We demonstrate that the representations our model learns for these images cluster naturally into more categories than exist manually assigned ground-truth labels, suggesting that the current labels may be overly coarse and may obscure important information about auroral morphology that may be useful in understanding the connection between auroral morphology and magnetospheric dynamics.

The rest of the paper is organized as follows: in Section III, we describe our model, training strategy, and results. In Section V, we conclude with a discussion of future work.

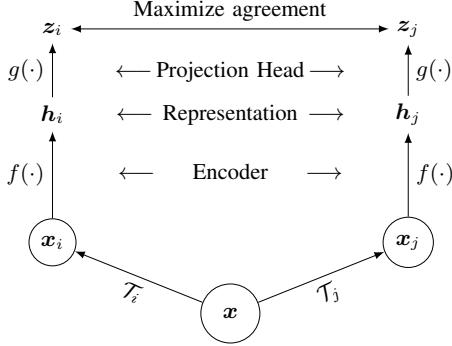


Fig. 2: A simple framework for contrastive learning of visual representations. Given a batch of images, for each sample \mathbf{x} , two distinct transformations are randomly applied (\mathcal{T}_i and \mathcal{T}_j) to obtain two correlated views of \mathbf{x} , a ‘positive pair’. A base encoder network $f(\cdot)$ and a projection head $g(\cdot)$ are then trained to identify the positive pairs among pairs constructed from all transformed samples in a minibatch. After training is completed, the encoder $f(\cdot)$ is applied to each \mathbf{x} to obtain a representation \mathbf{h} which can then be used for downstream tasks. Diagram adapted from [19].

III. METHODOLOGY

A. Data

The dataset used in this study consists of 5,824 images collected from various THEMIS all-sky imagers. This dataset was initially constructed for the study presented in [12], and is publicly available at http://tid.uio.no/plasma/oath/oath_v1.1_20181026.tgz. We refer to this dataset as the Oslo Auroral THEMIS (OATH) dataset.

Each image \mathbf{x} in the dataset was cropped by 15% to remove pixels corresponding to very low elevation angles and then scaled to enhance dim features by applying the following formula:

$$x_{ij} = \max\left(\min\left(\frac{x_{ij} - m_{\mathbf{x}}}{M_{\mathbf{x}}}, 1\right), 0\right), \quad (1)$$

where x_{ij} denotes the i, j -th pixel value, $m_{\mathbf{x}}$ denotes the 1st percentile brightness value, and $M_{\mathbf{x}}$ denotes the 99th percentile brightness of $\mathbf{x} - m_{\mathbf{x}}$. The resulting cropped and scaled images were then manually assigned a distinct label from the categories *Arc* (774 images), *Discrete* (1102 images), *Diffuse* (1400 images), *Cloudy* (817 images), *Moon* (585 images), *Clear/No aurora* (1082 images). These categories are described in [12] as follows:

- *Arc*: Showing one or multiple bands of aurora that stretch across the field-of-view; typically, the arcs have well-defined, sharp edges.
- *Diffuse*: Large patches of aurora, typically with fuzzy edges.
- *Discrete*: Auroral forms with well-defined, sharp edges, that are, however, not arc-like.
- *Cloudy*: Dominated by clouds or the dome of the imager is covered with snow.

- *Moon*: Dominated by light from the Moon.
- *Clear/No aurora*: Images which show a clear sky (stars and planets are clearly visible) without the appearance of aurora.

Examples from each class are displayed in Figure 1.

B. SimCLR

The Simple framework for Contrastive Learning of Representations (SimCLR) algorithm was introduced by Chen *et. al.* in [19]. As the name suggests, SimCLR is a relatively simple and yet highly effective algorithm for unsupervised representation learning. The SimCLR framework consists of four components, described below and diagrammed in Figure 2:

- A stochastic *data augmentation module*. Given a sample \mathbf{x} , the data augmentation module randomly applies two transformations $\mathcal{T}_i, \mathcal{T}_j$ to \mathbf{x} , resulting in two distinct views $\mathbf{x}_i, \mathbf{x}_j$ of \mathbf{x} that we consider a *positive pair* in the context of the contrastive loss function described below.
- A *base encoder*. The base encoder is used to extract an initial representation of a transformed input $\mathcal{T}_k(\mathbf{x})$. In principle, this could be any machine learning model; in most applications a residual neural network (ResNet) [23], sometimes pretrained on ImageNet, is used.
- A *projection head*. The projection head maps the output of the base encoder to the space where the contrastive loss function is to be applied. Again, in principle this could be any machine learning model; in practice the projection head is usually a shallow (one-layer) dense neural network.
- A *contrastive loss function* designed for the following contrastive prediction task: given a set of pairs of examples that includes a positive pair, identify the positive pair.

The contrastive loss function used here is the *normalized temperature-scaled cross-entropy loss* [24]–[26]. Letting $\text{sim}(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{y} / \|\mathbf{x}\| \|\mathbf{y}\|$ (cosine similarity), the loss for a positive pair of examples $\mathbf{x}_i, \mathbf{x}_j$ is given by

$$\ell(\mathbf{x}_i, \mathbf{x}_j) = -\log \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j) / \tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(\text{sim}(\mathbf{z}_k, \mathbf{z}_i) / \tau)}. \quad (2)$$

In Equation 2, $\mathbf{z}_i, \mathbf{z}_j$ are the outputs of the projection head for two views $\mathbf{x}_i, \mathbf{x}_j$ of an input \mathbf{x} . τ is a temperature parameter, $\mathbb{1}_{[k \neq i]}$ is an indicator function that evaluates to 1 if $i \neq k$, and the $2N$ examples represented in the sum in the denominator are obtained by using the $2(N - 1)$ augmented pairs in a minibatch as negative examples. This mitigates any need to sample negative pairs explicitly. The model is then trained to minimize

$$\frac{1}{2N} \sum_{k=1}^N [\ell(2k - 1, 2k) + \ell(2k, 2k - 1)]. \quad (3)$$

The choice of transformations sampled from is believed to be a critical component in learning successful representations using the SimCLR algorithm. We opt for simplicity, and use only two: a random resized cropping of the input image and a random flip over a horizontal axis. The latter choice is motivated by the view that the location of auroral forms in the image will be relevant to the ultimate goal of obtaining a better understanding of the magnetospheric dynamics involved; a vertical flip causes the model to judge as similar forms that occur on the left side of the image and forms that occur on the right, but a horizontal flip decorrelates the image pair while preserving the geographic location of the form.

In [19], it is noted that random color distortion, when paired with random cropping, is believed to be among the most important transformations to include in order to learn high-quality representations for image classification tasks. Our experience with this was mixed: incorporating random color distortions did improve classification performance somewhat beyond the results documented here, but it also degraded the quality of the representations learned for clustering and other downstream tasks. We therefore omitted random color distortions along with other commonly included transformations such as random rotations and perspective shifts.

In all of our experiments, we use a projection head consisting of two fully connected linear layers with an intermediate ReLU activation and without biases; that is, $g(\mathbf{h}) = W_2(\max(0, W_1(\mathbf{h})))$, where W_1 and W_2 denote weight matrices. As an encoder, we use a lightweight ResNet18 model that is randomly initialized. The temperature hyperparameter τ of the loss function (c.f. Eq. 2) was set to 0.5.

C. Training

We train the SimCLR model for 100 epochs, using the Adam optimizer [27] with a fixed learning rate of 0.0003 and no weight decay. We use a minibatch size of 128. All models were implemented using the open source machine learning library PyTorch [28] and trained using a single NVIDIA Titan Xp GPU.

D. Classification

After training, we fit an ℓ_2 -regularized logistic regression classifier to the obtained representations. We use 5-fold cross-validation to obtain average classification performance as measured by a range of metrics; c.f. Table I. For comparison, we also include results from fine-tuning an ImageNet-pretrained ResNet18. In order to insure maximum comparability between models we use the same folds for cross-validation as in [12]. Linear classifiers were fit using Scikit-Learn [29].

A confusion matrix for one of the held-out validation sets is presented in Figure 3. Confusion matrices for the other folds are very similar.

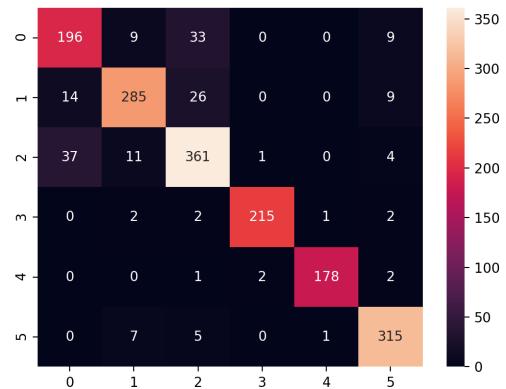


Fig. 3: **Confusion matrix for a single fold.** Rows are ground-truth labels, columns are predicted labels.

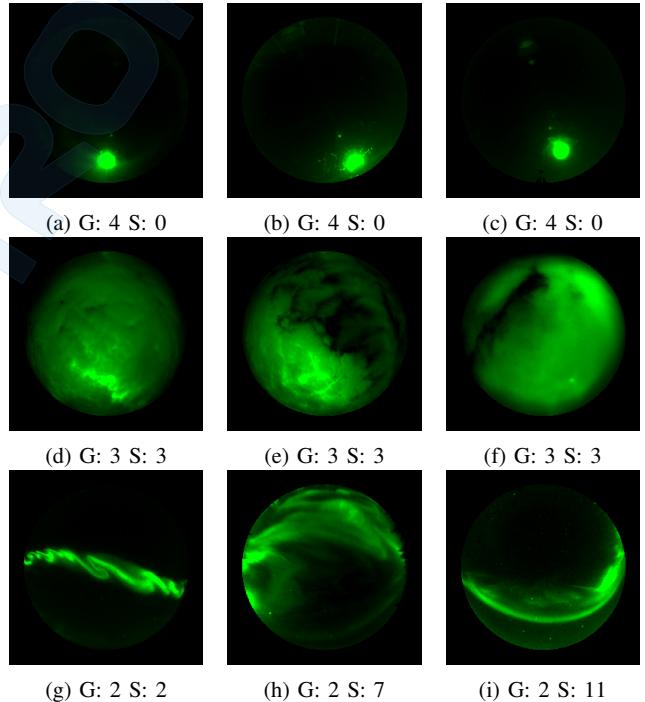


Fig. 4: A random sample of images from three ground-truth categories and their cluster labels. ‘G’ indicates the ground-truth label (0=‘Arc aurora’, 1=‘Diffuse aurora’, 2=‘Discrete aurora’, 3=‘Cloudy’, 4=‘Moon’, 5=‘Clear/No aurora’); ‘S’ the category assigned by the K-means algorithm to the learned representation of the image. In rows 1 and 2, the categories largely agree. In row 3, the ‘Discrete aurora’ category, which can be seen in the sample to contain very different auroral forms, is split over several distinct clusters.

Model	Params.	Top-1 Accuracy (%)	Std.	Precision	Recall	F1-Score
Ridge + Inception (pretrained) [12]	43m	81.7	0.1	83.0*	84.3*	83.6*
ResNet18 (pretrained, supervised)	12m	84.7	0.7	85.6	85.2	85.3
SimCLR	12m	90.9	0.7	91.6	91.5	91.5

TABLE I: **Auroral image classification results.** Data in row 1 is taken from [12]. The starred entries in row 1 were not reported in the cited paper, and instead were calculated from a single fold confusion matrix provided in the paper. Supervised ResNet18 results included for comparison. Reported results are averages from 5-fold cross-validation. To optimize comparability, the same folds are used for cross-validation as in [12].

E. Clustering

After training, the K -means clustering algorithm was applied to cluster the learned representations. By comparing average silhouette scores for $k \in \{3, \dots, 15\}$, we find evidence that the number of clusters present in the data is greater than the number of ground-truth labels currently used to classify the dataset. Specifically, we attain a maximum average silhouette score of 0.212 for $k = 12$. For comparison, using the ground-truth labels as cluster assignments yields an average silhouette score of 0.011. Significantly, while some clusters align closely with the ground-truth classifications (the classes *Cloudy* and *Moon*, for example), other less well-defined ground-truth classes (in particular, the *Discrete* class) are split among various clusters; c.f. Figure 4, 5. Moreover, random samples drawn from the learned clusters exhibit clear qualitative similarities; c.f. Figure 5.

IV. DISCUSSION

There are several distinctions to be made between the approach outlined here and the standard SimCLR model. As mentioned previously, we greatly reduce the number of transformations sampled from and nevertheless obtain high performance. We train on a small dataset using a lightweight model with 12m parameters, and improve on the current state-of-the-art, suggesting that while the conventional wisdom is that to learn effective representations unsupervised models should be more highly parameterized than their supervised counterparts, this may in fact not generally be the case.

In [30], it is noted that classification error rates of less than 10% are likely to be sufficient for most practical purposes; in particular, for automatic labeling of auroral images to take place on a large scale, thereby enabling statistical studies of auroral images not previously possible. The results presented in Section III-D, Table I, and Figure 3 demonstrate that the approach described here meets this criteria. Moreover, the model is relatively lightweight, using slightly less than 12m parameters and requiring less than 110MB of space to train to convergence. We obtain our state-of-the-art results in spite of the fact that our model does not rely on ImageNet pretraining, the usual approach when working with similarly sized datasets, thus removing another potential source of bias [18].

The results presented in Section III-E align with the general fact that the morphology of auroral forms is not well understood. The approach outlined here leads to representations that provide a much finer categorization of auroral forms, with clear qualitative distinctions between categories. This suggests

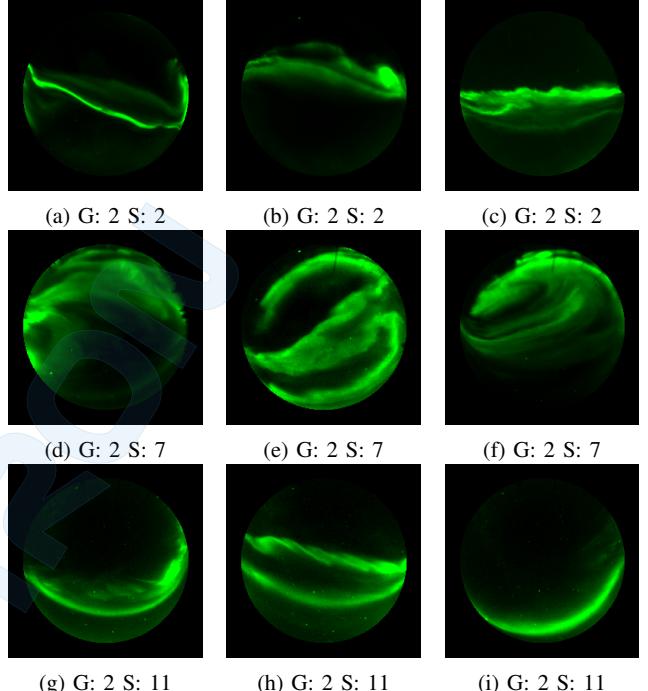


Fig. 5: A random sample of images from the ‘Discrete aurora’ category in specific cluster assignments. ‘G’ indicates the ground-truth label; ‘S’ the cluster label (c.f. Figure 4). The images in each row share characteristics that distinguish them from the images in the other rows.

that unsupervised learning may well have an important role to play in improving our understanding of the connection between what is viewed from the ground and the dynamics and coupling of Earth’s magnetosphere and upper atmosphere.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented a novel approach based on adapting the SimCLR model introduced in [19] to the practical challenge of auroral image classification. Our approach leads to state-of-the-art results as measured by a range of classification metrics, surpassing an important threshold for practical utility while requiring less than 25% of the parameters of recent benchmark models and without relying on ImageNet pretraining. Our approach provides preliminary evidence to suggest that current choices of ground-truth labels for auroral

images are likely overly coarse. We also demonstrate that the guidelines for learning effective representations put forth in [19] may be ideal for conventional object recognition tasks, but may require adjustment for data that does not share the same general characteristics. More research is needed to better understand the connection between the representations learned by our models and the magnetospheric processes involved. This will be a focus of future work.

REFERENCES

- [1] “NASA Science Mission Directorate (SMD) 2014 science plan,” <https://science.nasa.gov/about-us/science-strategy>, accessed 1 July 2021.
- [2] “2013 National Research Council decadal strategy for solar and space physics report, solar and space physics: A science for a technological society,” https://www.nap.edu/catalog.php?record_id=13060, accessed 1 July 2021.
- [3] S.-I. Akasofu, “The development of the auroral substorm,” *Planetary and Space Science*, vol. 12, no. 4, pp. 273–282, 1964.
- [4] ———, “Energy coupling between the solar wind and the magnetosphere,” *Space Science Reviews*, vol. 28, no. 2, pp. 121–190, 1981.
- [5] L. Clausen, S. E. Milan, and A. Grocott, “Thermospheric density perturbations in response to substorms,” *Journal of Geophysical Research: Space Physics*, vol. 119, no. 6, pp. 4441–4455, 2014.
- [6] P. Newell, T. Sotirelis, and S. Wing, “Diffuse, monoenergetic, and broadband aurora: The global precipitation budget,” *Journal of Geophysical Research: Space Physics*, vol. 114, no. A9, 2009.
- [7] M. Syrjäsu and E. Donovan, “Analysis of auroral images: detection and tracking,” *Geophysica*, vol. 38, no. 1-2, pp. 3–14, 2002.
- [8] ———, “Diurnal auroral occurrence statistics obtained via machine vision,” in *Annales Geophysicae*, vol. 22, no. 4. Copernicus GmbH, 2004, pp. 1103–1113.
- [9] M. Syrjäsu, E. Donovan, X. Qin, and Y. Yang, “Automatic classification of auroral images in substorm studies,” in *8th International Conference on Substorms (ICS8)*, 2007, pp. 309–313.
- [10] J. Rao, N. Partamies, O. Amariti, M. Syrjäsu, and K. E. van de Sande, “Automatic auroral detection in color all-sky camera images,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 7, no. 12, pp. 4717–4725, 2014.
- [11] Q. Yang, J. Liang, Z. Hu, and H. Zhao, “Auroral sequence representation and classification using hidden markov models,” *IEEE transactions on geoscience and remote sensing*, vol. 50, no. 12, pp. 5049–5060, 2012.
- [12] L. B. Clausen and H. Nickisch, “Automatic classification of auroral images from the oslo auroral themis (oath) data set using machine learning,” *Journal of Geophysical Research: Space Physics*, vol. 123, no. 7, pp. 5640–5647, 2018.
- [13] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [14] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *CVPR09*, 2009.
- [15] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, “Cnn features off-the-shelf: an astounding baseline for recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2014, pp. 806–813.
- [16] A. Kvammen, K. Wickstrøm, D. McKay, and N. Partamies, “Auroral image classification with deep neural networks,” *Journal of Geophysical Research: Space Physics*, vol. 125, no. 10, p. e2020JA027808, 2020.
- [17] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, “Cnn features off-the-shelf: An astounding baseline for recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2014.
- [18] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F. A. Wichmann, and W. Brendel, “Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness,” *arXiv preprint arXiv:1811.12231*, 2018.
- [19] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [20] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, “Unsupervised learning of visual features by contrasting cluster assignments,” *arXiv preprint arXiv:2006.09882*, 2020.
- [21] Z. Dai, B. Cai, Y. Lin, and J. Chen, “Up-detr: Unsupervised pre-training for object detection with transformers,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 1601–1610.
- [22] E. Donovan, S. Mende, B. Jackel, H. Frey, M. Syrjäsu, I. Voronkov, T. Trondsen, L. Peticolas, V. Angelopoulos, S. Harris *et al.*, “The themis all-sky imaging array—system design and initial results from the prototype imager,” *Journal of Atmospheric and Solar-Terrestrial Physics*, vol. 68, no. 13, pp. 1472–1487, 2006.
- [23] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [24] K. Sohn, “Improved deep metric learning with multi-class n-pair loss objective,” in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, 2016, pp. 1857–1865.
- [25] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, “Unsupervised feature learning via non-parametric instance discrimination,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3733–3742.
- [26] A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018.
- [27] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [28] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [29] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [30] M. Syrjäsu and N. Partamies, “Numeric image features for detection of aurora,” *IEEE Geoscience and Remote Sensing Letters*, vol. 9, no. 2, pp. 176–179, 2011.