

# COLLABORATIVE LEARNING OF SEMI-SUPERVISED CLUSTERING AND CLASSIFICATION FOR LABELING UNCURATED DATA

Sara Mousavi\* Dylan Lee\* Tatianna Griffin<sup>†</sup> Dawnie Steadman<sup>†</sup> Audris Mockus\*

\* Department of Electrical Engineering and Computer Science

<sup>†</sup>Department of Anthropology

The University of Tennessee Knoxville, USA

## ABSTRACT

Domain-specific image collections present potential value in various areas of science and business but are often not curated nor have any way to readily extract relevant content. To employ contemporary supervised image analysis methods on such image data, they must first be cleaned and organized, and then manually labeled for the nomenclature employed in the specific domain, which is a time consuming and expensive endeavor. To address this issue, we designed and implemented the Plud system. Plud provides an iterative semi-supervised workflow to minimize the effort spent by an expert and handles realistic large collections of images. We believe it can support labeling datasets regardless of their size and type. Plud is an iterative sequence of unsupervised clustering, human assistance, and supervised classification. With each iteration 1) the labeled dataset grows, 2) the generality of the classification method and its accuracy increases, and 3) manual effort is reduced. We evaluated the effectiveness of our system, by applying it on over a million images documenting human decomposition. In our experiment comparing manual labeling with labeling conducted with the support of Plud, we found that it reduces the time needed to label data and produces highly accurate models for this new domain.

**Index Terms**— image labeling, image clustering, image classification, image feature embedding, convolutional neural networks

## 1 Introduction

Most published image analysis models tend to assume clean, organized image collections with ample amount of samples for training purposes. Use of these advanced models in domains that have not yet embraced such approaches, or for image collections that have not undergone an expansive organization, cleaning, and other type of curation, is not, unfortunately, possible. The effort we spent to overcome this problem in a collection of images documenting human decomposition led us to build the Plud system<sup>1</sup> to support such tasks.

In the workflow that Plud provides, users start with a large and relatively unstructured set of images and produce highly accurate domain-specific models that can be used to clean, structure, and label the collection. That, in essence, unlocks the full potential of the content hidden in these images to solve research and practical problems. The main components that enables Plud to achieve its goals are: 1) the use of unsupervised clustering approaches using embeddings from an arbitrary domain to explore and design an initial grouping of images, 2) a user interface that supports rapid labeling of large batches of samples, and 3) iterations that continuously increase the accuracy of the model while reducing the labeling effort of the user.

Specifically, the iterative sequence of clustering, labeling, and classification gradually introduce new unseen data in each iteration. The resulted labels from each iteration are used as input data for a classification model in the next iteration. High confidence results from the classifier are used to boost its learning, while low confidence decisions route images to the clustering and manual labeling component in order to be corrected. Repeating the iterations results in more labeled data and more accuracy for the classifier, and hence less need for manual labeling in consequent iterations.

Using Plud, we classified over one million images depicting the human decomposition process and evaluated the performance of our classification on a subset of 5555 randomly selected and manually labeled images. The results show we were able to classify these images with top-1 and top-3 F-score of 79.89 and 93.84 respectively.

Using unsupervised methods one can cluster image data using their numerical embeddings [1, 2, 3, 4] in groups that share similar characteristics. However, there is no certainty that such methods will be able to recognize any relevant classes for a new domain without manual evaluation, which is time consuming.

Other works on semi-supervised and weakly supervised learning methods [5, 6, 7, 8, 9] target speeding up the data labeling process and other computer vision tasks. These approaches start with a small set of labeled data and a large set of unlabeled data. The labeled data then guides the learn-

<sup>1</sup>Link to the repository: [https://github.com/saramsv/body\\_part\\_classifier](https://github.com/saramsv/body_part_classifier)

ing process to make the models more generalizable to the remaining unlabeled data. In these works, transfer learning and fine-tuning are often used to extend what is learned from one dataset to other datasets [10, 11]. Unfortunately, labeled data that can represent the new dataset is needed for transfer learning and fine-tuning to perform well. This is especially the case when the new dataset has a completely different set of classes.

Ways to find samples that are hard to learn was considered in [12], and we incorporate similar techniques in Plud. A workflow for labeling by Philbrick et al. [13, 14] supports annotating medical images involving an iterative approach of manual labeling and classification. However, their system works well only for small medical datasets.

The rest of the paper is as follows. Detailed information about our proposed system is provided in Section 2. Section 3 includes an introduction to the dataset and the conducted experiments and their results. We finally conclude the paper in Section 4.

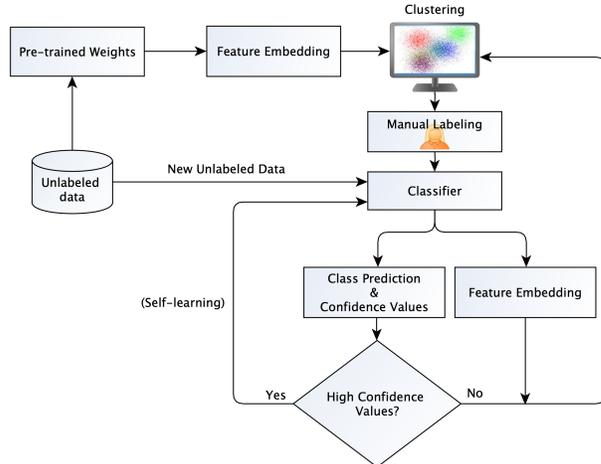
## 2 Method

Figure 1 shows an overview of Plud. The basic idea of Plud is to combine unsupervised clustering, and supervised classification methods in an iterative workflow involving a human expert, with a user interface tailored for enabling rapid labeling of a large number of images. In each iteration, the uncertainty of the classifier is used to suggest additional data to be clustered and then manually labeled. Images that are considered similar are presented in large batches to leverage the power of the human visual system to detect outliers. The domain expert produces names/labels for each group based on the dominant class in the cluster. The resulting labeled data are then used for training and fine-tuning the classifier in the next iteration.

In the following, details about data preparation, clustering, the labeling interface, and the classification step are provided.

### 2.1 Data Preparation

The performance and generalizability of classifiers depends on the quality and the size of their training data. The more diverse and representative the training data is, the more generalizable the models are likely to be. In the case of a temporal dataset, such as images documenting human decomposition in which the subjects' appearance changes over time, it is important to sample the training data in such a way that includes the dataset's characteristic. In our case, for example, images of all possible decomposition phases should be included in the training data. To ensure the inclusion of changes over time on the human bodies, we randomly selected a small number of subjects and then selected all images taken from them over time, from fresh to decay, rather than randomly selecting images from the entire set which might include more subjects but might also exclude some decomposition phases.



**Fig. 1:** Shows the overall architecture of Plud. Unlabeled data are mapped to numerical feature embeddings and then clustered together. The resulted clusters are displayed to the domain expert for manual labeling through the Plud interface. The labeled data is used for training a classifier by which, for a new set of unlabeled data, labels and feature embeddings are generated. The iteration repeats for images with low confidence predictions. High confidence predictions are fed back to the classifier to enable self-learning.

### 2.2 Clustering

In order to assist and speed up the manual labeling process and enable mass labeling, we built a web-based interface that clusters images so that they can be viewed and evaluated by users. To build the clusters, the unlabeled data are mapped to feature embeddings using a convolutional neural network pre-trained on Imagenet [15]. In the first iteration that we had no labeled data, we used a ResNet-based model. We fed in the images to the network and extracted the output of the convolutional layers as the feature embeddings and then clustered them. For subsequent iterations, we used the feature embeddings generated by our trained classifier model. We used various clustering methods such as agglomerative clustering [16], spectral clustering [17] and KMeans [18] to group similar images together. From a manual examination of the clusters we concluded that KMeans outperformed the rest of the methods for our dataset.

The resulted clusters are displayed to a user for evaluation and cleaning through a web interface. The interface allows the user to see all images assigned to one cluster in a single web page. Each image can be selected by a single click if they are misclustered. The rest of the cluster is labeled with the label that matched the majority of the images in the cluster.

### 2.3 Classification

We use the labeled data resulting from the clustering-based approach to train a classifier model. The output of our classifier

for each input is a feature embedding which is resulted from the convolutional layers only, a label and a confidence value for the given label. The initial training data for the classifier is obtained from the labels given by the expert to the initial clusters. During subsequent iterations, the classifier is used to make predictions on the batches of unlabeled data, and the predictions are ordered by the confidence level the classifier assigns to each prediction. High-confidence predictions are used to expand the training data using the predicted labels and low confidence predictions are assigned for clustering and manual labeling. The threshold for considering high or low confidence values is determined by manually exploring the predictions.

Adding images with high confidence predicted labels to the classifier’s training data improve its accuracy on those type of images. Manual labeling of the images with low confidence predictions is still necessary to accurately predict the class for the instances where the certainty of the classifier was low. In other words, the low confidence values indicate the types of data the classifier has not yet fully learned. To expose the classifier to such data, we use the feature embedding outputted from the classifier for these images to cluster them. The clusters are then displayed to a domain expert for labeling and verification. Note that these embeddings are better representations of the images than those from the initial Imagenet-trained model as they have been partially trained on the data from the domain.

## 3 Evaluation

We have tested our method on a dataset of images capturing the process of human decomposition. In the following section we provide more details about the dataset, the conducted experiments and the end results.

### 3.1 Dataset

Images tracking human decomposition process are of great value for conducting research. This is due to two main reasons. First, when using images instead of actual decaying bodies, studying the human decomposition process is no longer limited to only the period of decomposition itself. Second, it allows research to focus on particular factors, such as mold formation or mummification, and correlate such factors to both environmental and individual characteristics (e.g., age or weight).

The image collection consists of photos taken of decomposing humans donated to the Forensic Anthropology Center of our university. These donors are placed at the Anthropology Research Facility (“Body Farm”) where the different stages of decomposition are studied and recorded through photography. An identity number is assigned to each donor and written on a wooden stake and placed next to them. The pictures are taken periodically from various angles of the bodies and capture different body parts to show the stage of body decomposition

at any given time. The collection has over one million images and images from following categories: *Arm, Hand, Foot, Legs, Full Body, Head, Backside, Torso, Stake, Plastic*.

For this work, we selected and labeled 5555 images capturing all images taken from 3 subjects as our ground truth data for testing Plud. The performance metrics (average precision, average recall and F-score) reported in this paper are resulted using this test data.

## 3.2 Results

In this section, we evaluate the efficiency gained from our labeling interface in addition to the affects of data sampling and the size of training data on the labeling performance.

### 3.2.1 Labeling Interface Performance

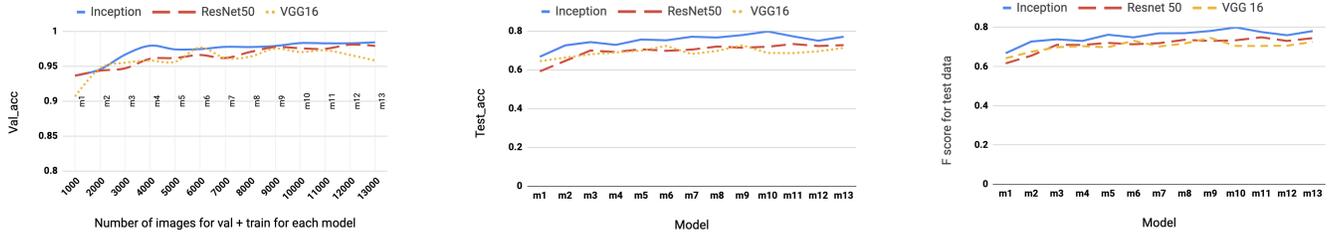
In order to examine how our clustering-based web interface affects the speed of labeling, we randomly selected 150 images and asked a user to label them with and without using our interface. We then repeated the same process for 300 images to see if the labeling time linearly increases when we double the number of images. In order to have a fair comparison, we designed a similar interface for the user with and without clusters. For the basic case that does not involve clustering, image thumbnails are generated on a single web page where the user can select the desired images. Using this interface, the user selects images belonging to the same class in each iteration until all images are labeled with appropriate classes. In the cluster-based interface of Plud, images are displayed based on the cluster they belong to. Thus the user only needs to select the miss-clustered images and label the cluster based on the majority of the images within it. The result of this experiment is shown in Table 2. The result shows that our interface drastically reduces the labeling time. It also shows that a linear increase in the number of images non-linearly increases the time of labeling. The result also indicates that labeling smaller batches of data is faster. We believe that this is due to the visual confusion caused by scrolling through large number of images.

### 3.2.2 Data Sampling Strategy

As mentioned in Section 2.1, it is important to make sure that the train data is diverse and represents the structure of the dataset. We designed an experiment to test this hypothesis by having a set of 5000 randomly selected images from the entire dataset and 5000 images from multiple randomly selected subjects while making sure it includes all stages of decomposition. We trained our classifier on the two sets and tested it on our test set. The results shown in Table 3 indicate that the classifier trained on the structure-aware multi-subject data performs with 76.71% accuracy on the test set where as the accuracy is 74.51% for the classifier trained on randomly selected data.

### 3.2.3 Classifier Performance

We analyzed the effect of adding data in each iteration on the performance of the classifier to ensure that in each iteration we



**Fig. 2:** The performance of models trained on various number of images is shown. Test\_acc and F score are calculated using the manually labeled test data mentioned in Section 3.1.  $\{m_1, m_2, \dots, m_{13}\}$  are the models that are refined with more data through 13 iterations. Although it is expected to have better performance as the amount of training data increases, it is important to note that, the more accuracy in each iteration results in less effort for the user in the labeling process as shown in Table 2.

Model		Precision of Classes										AP
		Arm	Hand	Foot	Legs	Full Body	Head	Backside	Torso	Stake	Plastic	
Inception	Top 1	45.73	85.60	93.72	60.52	92.69	94.33	68.25	87.22	96.30	73.61	<b>79.80</b>
	Top 3	80.23	96.86	97.75	86.96	97.53	98.29	89.57	97.20	98.87	88.88	<b>93.21</b>
Model		Recall of Classes										AR
		Arm	Hand	Foot	Legs	Full Body	Head	Backside	Torso	Stake	Plastic	
Inception	Top 1	53.88	67.36	62.34	97.60	77.21	94.91	55.84	91.97	98.86	1	<b>80.00</b>
	Top 3	92.69	87.63	90.66	99.57	96.34	99.75	81.81	96.27	1.	1.	<b>94.47</b>

**Table 1:** Shows precision and recall for our Inception-based classifier on the test data.

	150 images	300 images
Basic labeling	13m.22s	43m.52s
Unsupervised assisted labeling	4m.41s	11m.34s

**Table 2:** In the Basic method, the user sees all images at once and labels them. In the second approach, the images are first grouped using our clustering interface and the user labels the groups and corrects the misclustered images. The time for 300 images shows that the increase in the number of images non-linearly adds to the labeling time.

have a more accurate model and thus less effort on the domain expert for manual labeling. We used three different models to find the model with highest accuracy. We checked the performance of VGG16, ResNet50 and Inception-based classifiers trained on various amounts of data,  $\{1000, 2000, \dots, 13000\}$ . In each iteration, a new set of 1000 unlabeled images were fed to the classifier. High confidence predictions were used as new labels and images with low confidence predictions were labeled through the clustering interface. The new 1000 labeled images were added to the training data.

When the classifier is trained, its performance on a test set of 5555 manually labeled images was examined. The vall\_acc, test\_acc and F-score are shown in Figure 2. Increasing the size of the training data in all three classifiers resulted in an overall increase in their performance on our test data. It is important to note that although this result is expected, in the context of Plud that means less manual effort on labeling for

	accuracy
Random sampling	74.51
Structure-aware multi random subject sampling	76.71

**Table 3:** Shows the performance of an Inception based classifier when trained on randomly selected images vs random multi-subjects' data.

the domain expert since less misclustered images need to be manually selected each time.

Based on the results shown in Figure 2, we obtained the best accuracy when using an Inception-based classifier which is then used to calculate precision and recall for test data shown in Table 1.

## 4 Conclusion

We proposed Plud as a system for speeding up the labeling of uncured large collections of images in cases where no labeled data pre-exists and the types of images, and relevant image classes are difficult to assess or are not available in advance. For evaluation, we used Plud to label images depicting human decomposition. Our evaluation shows that Plud accelerates and facilitates labeling of large amounts of data. We were able to easily label 13192 and 5555 train and test images respectively using the iteration provided in Plud. Our trained classifier on these labeled data was then used to automatically label the remaining 1, 148, 511 images.

## 5 References

- [1] Joris Guérin, Olivier Gibaru, Stéphane Thiery, and Eric Nyiri, “Cnn features are also great at unsupervised classification,” *arXiv preprint arXiv:1707.01700*, 2017.
- [2] Jianwei Yang, Devi Parikh, and Dhruv Batra, “Joint unsupervised learning of deep representations and image clusters,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 5147–5156.
- [3] Sara Mousavi, Dylan Lee, Tatianna Griffin, Dawnie Steadman, and Audris Mockus, “An analytical workflow for clustering forensic images,” *arXiv preprint arXiv:2001.05845*, 2019.
- [4] Xiaosong Wang, Le Lu, Hoo-Chang Shin, Lauren Kim, Mohammadhadi Bagheri, Isabella Noguees, Jianhua Yao, and Ronald M Summers, “Unsupervised joint mining of deep features and image labels for large-scale radiology image categorization and scene recognition,” in *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2017, pp. 998–1007.
- [5] Inmaculada Dópido, Jun Li, Prashanth Reddy Marpu, Antonio Plaza, José M Bioucas Dias, and Jon Atli Benediktsson, “Semisupervised self-learning for hyperspectral image classification,” *IEEE transactions on geoscience and remote sensing*, vol. 51, no. 7, pp. 4032–4044, 2013.
- [6] Chuck Rosenberg, Martial Hebert, and Henry Schneiderman, “Semi-supervised self-training of object detection models.,” *WACV/MOTION*, vol. 2, 2005.
- [7] David Yarowsky, “Unsupervised word sense disambiguation rivaling supervised methods,” in *33rd annual meeting of the association for computational linguistics*, 1995, pp. 189–196.
- [8] Hye-Woo Lee, Noo-ri Kim, and Jee-Hyong Lee, “Deep neural network self-training based on unsupervised learning and dropout,” *International Journal of Fuzzy Logic and Intelligent Systems*, vol. 17, no. 1, pp. 1–9, 2017.
- [9] Luigi Grimaudo, Marco Mellia, Elena Baralis, and Ram Keralapura, “Select: Self-learning classifier for internet traffic,” *IEEE Transactions on Network and Service Management*, vol. 11, no. 2, pp. 144–157, 2014.
- [10] Mohammad Raji, Alok Hota, Robert Sisneros, Peter Messmer, and Jian Huang, “Photo-guided exploration of volume data features,” *arXiv preprint arXiv:1710.06815*, 2017.
- [11] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten, “Exploring the limits of weakly supervised pretraining,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 181–196.
- [12] Jiabin Xue, Jiqing Han, Tieran Zheng, Jiaying Guo, and Boyong Wu, “Hard sample mining for the improved re-training of automatic speech recognition,” *arXiv preprint arXiv:1904.08031*, 2019.
- [13] Kenneth A Philbrick, Alexander D Weston, Zeynettin Akkus, Timothy L Kline, Panagiotis Korfiatis, Tomas Sakinis, Petro Kostandy, Arunnit Boonrod, Atefeh Zeinoddini, Naoki Takahashi, et al., “Ril-contour: a medical imaging dataset annotation tool for and with deep learning,” *Journal of digital imaging*, pp. 1–11, 2019.
- [14] Yanchun Ma, Yongjian Liu, Qing Xie, and Lin Li, “Cnn-feature based automatic image annotation method,” *Multimedia Tools and Applications*, vol. 78, no. 3, pp. 3767–3780, 2019.
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [16] Daniel Müllner, “Modern hierarchical, agglomerative clustering algorithms,” *arXiv preprint arXiv:1109.2378*, 2011.
- [17] X Yu Stella and Jianbo Shi, “Multiclass spectral clustering,” in *null*. IEEE, 2003, p. 313.
- [18] K Krishna and Narasimha M Murty, “Genetic k-means algorithm,” *IEEE Transactions on Systems Man And Cybernetics-Part B: Cybernetics*, vol. 29, no. 3, pp. 433–439, 1999.