

The iMet Collection 2019 Challenge Dataset

Chenyang Zhang¹, Christine Kaeser-Chen¹, Grace Vesom⁵, Jennie Choi³, Maria Kessler³, and Serge Belongie^{1,3}

¹Google Research

²The Metropolitan Museum of Art

³Cornell University and Cornell Tech

⁵Arraiy, Inc.

Abstract

Existing computer vision technologies in artwork recognition focus mainly on instance retrieval or coarse-grained attribute classification. In this work, we present a novel dataset for fine-grained artwork attribute recognition. The images in the dataset are professional photographs of classic artworks from the Metropolitan Museum of Art, and annotations are curated and verified by world-class museum experts. In addition, we also present the iMet Collection 2019 Challenge as part of the FGVC6 workshop. Through the competition, we aim to spur the enthusiasm of the fine-grained visual recognition research community and advance the state-of-the-art in digital curation of museum collections.

1. Introduction

Fine grained visual categorization (FGVC) has emerged as a popular research area in computer vision in recent years. With the rapid progress of deep neural networks, computer vision algorithms can now capture powerful representations for complex semantics in domains such as fashion and biodiversity. Large-scale datasets such as COCO [6] and ImageNet [2] have played important roles in helping advance state-of-the-art algorithms for coarse-grained category recognition, and the emergence of FGVC datasets provide a complementary role for a wide variety of subcategories.

Previous works on artwork recognition focus mainly on instance retrieval [1, 4], mid-level attributes such as color [9], and person detection [3]. Other datasets feature descriptive attributes such as style, theme and mood [5, 8].

In this work, we present a novel dataset, *iMet Collection 2019*, focusing on fine-grained artwork attribute recogni-

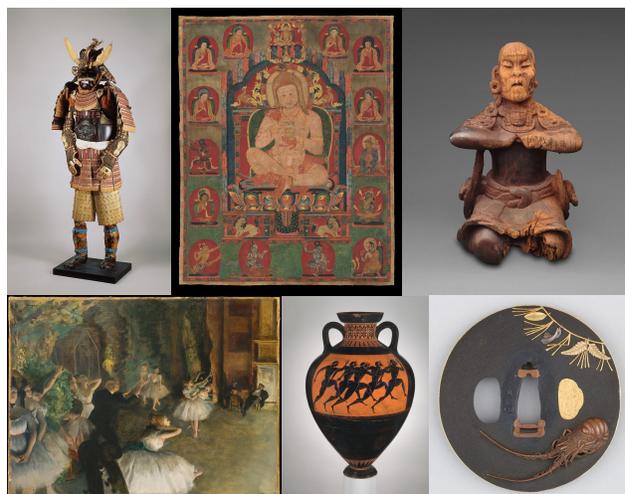


Figure 1. Sample images from the iMet Collection dataset.

tion. Compared to previous artwork datasets, the proposed dataset features the following two characteristics:

- **Fine-grained** The proposed dataset contains diverse attributes from a specialized domain.
- **Research-grade** Museum experts curated and verified attribute labels to ensure high quality.

2. The iMet Collection Dataset

The Metropolitan Museum of Art in New York (The Met) has a diverse collection of artworks from across the globe spanning 5,000 years. The artworks in The Met’s collection include paintings, musical instruments, prints, arms and armor, clothing, sculpture, furniture, metalwork, and more.

In 2017 The Met launched its Open Access Program making over 375,000 artwork images and data in its collection available under Creative Commons CC0 license. The

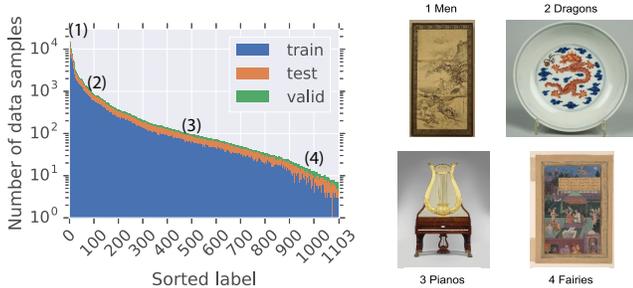


Figure 2. Left: distribution of images per attribute label. Right: representative images for 4 sampled attributes.

Open Access data includes object type, artist, title, period, creation date, medium, culture, dimensions, and geographic information. The Met recently added subject keyword tags to their Open Access Program to improve discoverability and increase access to the collection. Together the data and the subject keyword tags can also be used for content suggestion and a better means to navigate the collection.

To scale the coverage of the subject keyword tags to all artwork images, we might need to leverage more automated annotation algorithms. To that end, we release the iMet Collection dataset consisting of images with groundtruth keyword labels, which could facilitate the research in computer vision algorithms for keyword annotation as well as other applications.

2.1. Images

All images from our dataset are supplied by The Met and are all publicly available under CC0 license.¹ While there may be more than one image available for each *object* (artwork instance), we only retain the ones for “main-display.”² Each image is resized such that the shorter dimension is 300 pixels and is encoded in PNG format. Sample images are shown in Figure 1.

2.2. Annotations

The Met’s subject keyword tags cover religious subjects, historical events, plants and animals. The tag set provides a new perspective into the collection, allowing users to search common topics across time.

The annotations of the dataset come from two sources: (1) Subject Matter Experts (SME) and (2) vendor annotators.

SME derived attributes are provided at curation time by the SME, and may contain single or multiple labels. These labels are drawn from several predefined domains such as *Culture, Medium, Origin, etc.* For this competition, we only include the *Culture* domain for simplicity.

¹The Met Online Collection: <https://www.metmuseum.org/art/collection>

²*Main-display* refers to images capturing the most comprehensive view, often frontal, of the object.

Vendor-labeled attributes were added by a professional outsourced organization from a taxonomy provided by The Met. The annotators were able to view the museum’s online collection pages and were advised to avoid annotating labels already present. Specifically, the vendor was advised to annotate labels which can be visually inferred (e.g. “Shellfish” in Figure 1, bottom-right sword guard). These annotations can be potentially more intuitive and visually plausible.

We have included 1103 different attributes from both culture-related attributes from SME derived attributes and vendor-sourced attributes in the final dataset.

2.3. Data Splits and Sizes

We divide the dataset into three splits: *training*, *validation*, and *test* subsets. While the training and validation splits were released during the competition, we retained the test split for final ranking and will release it afterwards.

Since each data sample is annotated with 1 or more attributes, we first randomly pick a “pivot” attribute label for each sample. Then we split the dataset using the “pivot” attribute labels in the following manner:

1. We discard samples whose pivot attributes have fewer than three samples.
2. If an pivot attribute has fewer than 10 samples, we assign one to each of validation and test set, and the rest to training set.
3. Otherwise, we assign 70%, 5%, and 25% of them to training, validation, and test, respectively.

Finally, we reattach the non-pivot labels back to each assigned sample and discard labels which do not exist in all three splits. In the end, our dataset contains 1103 attribute labels from 155,531 samples, among which we have 109,274 samples for training, 7443 for validation and 38,814 for testing.

2.4. Data Challenges

The iMet Collection 2019 Dataset has a long-tailed distribution (Figure 2) over the attribute label set. The three most popular attributes are “Men”, “Women”, and “French” (which contain 19,974, 14,283, and 13,549 training samples, respectively), and the least popular tag is “Afgan”, having only one training sample. Another challenge is the unbalanced numbers of attributes per sample. As shown in Figure 3, the number of attributes per sample ranges between 1 and 11, while most of the samples have 2 to 4 attributes.

2.5. Metric

We evaluate the performances of our dataset following the multi-label classification manner and employ the F_β [7] score:

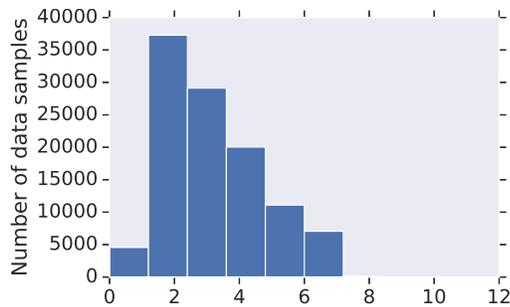


Figure 3. Histogram of per-sample attribute numbers over the training set. Most samples have 2 to 4 attributes while one of the training sample has 11 attributes.

$$F_{\beta} = \frac{(1 + \beta^2)TP}{(1 + \beta^2)TP + \beta^2FN + FP} \quad (1)$$

where TP, FN, and FP stand for numbers of *true positives*, *false negatives* and *false positives*, respectively. In the challenge, we set $\beta = 2$ to trade for more true positives with more tolerance to false positives.

3. The iMet Collection Challenge 2019

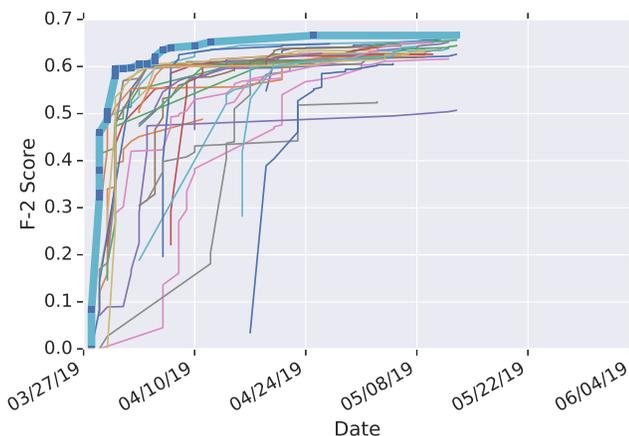


Figure 4. Challenge leader-board scores up to date. Each colored thin line represents a team with their submissions evaluated on validation set. The cyan thick curve indicates the leading score over time among all teams. We only plot 32 teams who have the longest date spans.

The iMet Collection Challenge 2019 is conducted through Kaggle as part of FGVC6 workshop at CVPR19. At the time of paper submission, over **500** competitors from **435** teams have submitted their results. We plot the public leaderboard progress over time in Figure 4. The competition is the first kernel-only Computer Vision challenge ever hosted on Kaggle, where the participants can only run their model within limited resources and are required to open-source their model (kernel) at the end of competition. We hope that this effort can further benefit the research community by increasing algorithm transparency and resource fairness.

4. Conclusions

Artwork attribute recognition is a novel area of research in fine-grained visual classification (FGVC) domain. In this work, we propose a novel multi-attribute image dataset to enrich the assets of this research area. The iMet Collection 2019 dataset is the first high-quality artwork image dataset with research-grade attribute labels curated or verified by world-class museum experts.

As the attribute annotation work is an ongoing effort in the Met Open Access program, we plan to extend the iMet Collection challenge dataset with richer attributes covering more semantic groups, such as medium and size. In addition, we also plan to support more computer vision applications such as instance detection and segmentation, caption generation, and cross-modality retrieval in the future competitions.

References

- [1] Elliot J Crowley and Andrew Zisserman. In search of art. In *European Conference on Computer Vision*, pages 54–70. Springer, 2014.
- [2] 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*, pages 248–255. Ieee, 2009.
- [3] Shiry Ginosar, Daniel Haas, Timothy Brown, and Jitendra Malik. Detecting people in cubist art. In *European Conference on Computer Vision*, pages 101–116. Springer, 2014.
- [4] David Gorisse, Matthieu Cord, Michel Jordan, Sylvie Philipp-Foliguet, and Frédéric Precioso. 3d content-based retrieval in artwork databases. In *2007 3DTV Conference*, pages 1–4. IEEE, 2007.
- [5] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, et al. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://github.com/openimages>*, 2:3, 2017.
- [6] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [7] David Martin Powers. Evaluation: from precision, recall and f-measure to roc, informedness, markedness and correlation. 2011.
- [8] Michael J Wilber, Chen Fang, Hailin Jin, Aaron Hertzmann, John Collomosse, and Serge Belongie. Bam! the behance artistic media dataset for recognition beyond photography. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1202–1211, 2017.
- [9] Marchenko Yelizaveta, Chua Tat-Seng, and Aristarkhova Irina. Analysis and retrieval of paintings using artistic color concepts. In *2005 IEEE International Conference on Multi-media and Expo*, pages 1246–1249. IEEE, 2005.