



Semi-supervised Human Pose Estimation in Art-historical Images

Matthias Springstein

matthias.springstein@tib.eu

TIB – Leibniz Information Centre for Science and
Technology
Hanover, Germany

Christian Althaus

christian.althaus@gmx.eu

TIB – Leibniz Information Centre for Science and
Technology &
Institute of Data Science, Leibniz University Hannover
Hanover, Germany

Stefanie Schneider

stefanie.schneider@itg.uni-muenchen.de

Ludwig Maximilian University of Munich
Munich, Germany

Ralph Ewerth

ralph.ewerth@tib.eu

TIB – Leibniz Information Centre for Science and
Technology &
L3S Research Center, Leibniz University Hannover
Hanover, Germany

ABSTRACT

Gesture as ‘language’ of non-verbal communication has been theoretically established since the 17th century. However, its relevance for the visual arts has been expressed only sporadically. This may be primarily due to the sheer overwhelming amount of data that traditionally had to be processed by hand. With the steady progress of digitization, though, a growing number of historical artifacts have been indexed and made available to the public, creating a need for automatic retrieval of art-historical motifs with similar body constellations or poses. Since the domain of art differs significantly from existing real-world data sets for human pose estimation due to its style variance, this presents new challenges. In this paper, we propose a novel approach to estimate human poses in art-historical images. In contrast to previous work that attempts to bridge the domain gap with pre-trained models or through style transfer, we suggest semi-supervised learning for both object and keypoint detection. Furthermore, we introduce a novel domain-specific art data set that includes both bounding box and keypoint annotations of human figures. Our approach achieves significantly better results than methods that use pre-trained models or style transfer.

CCS CONCEPTS

• **Computing methodologies** → **Object detection**; • **Information systems** → *Image search*; • **Theory of computation** → *Semi-supervised learning*; • **Applied computing** → **Arts and humanities**.

KEYWORDS

human pose estimation, semi-supervised learning, style transfer, art history

ACM Reference Format:

Matthias Springstein, Stefanie Schneider, Christian Althaus, and Ralph Ewerth. 2022. Semi-supervised Human Pose Estimation in Art-historical Images. In *Proceedings of the 30th ACM International Conference on Multimedia (MM '22)*.



This work is licensed under a Creative Commons Attribution International 4.0 License.

MM '22, October 10–14, 2022, Lisboa, Portugal

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9203-7/22/10.

<https://doi.org/10.1145/3503161.3548371>

'22), Oct. 10–14, 2022, Lisboa, Portugal. ACM, New York, NY, USA, 10 pages.
<https://doi.org/10.1145/3503161.3548371>

1 INTRODUCTION

As ‘language’ of non-verbal communication, gesture has been theoretically established since the 17th century [23]. Its relevance for the visual arts, however, has so far been expressed at most sporadically [1]: e.g., symbolically-performatively on the basis of the medieval law-book manuscript of the Heidelberg *Sachsenspiegel* [45], as the antiquity-receiving ‘Pathosformel’ [48, 49], or as a status identifier exemplified in Roman sculpture [5]. This selectivity may be primarily due to the sheer overwhelming amount of data that traditionally had to be processed manually. Driven by the steady progress of digitization, though, an increasing quantity of historical artifacts has been indexed and made freely available to the public online in recent decades. As a result, art historians can draw on ever larger collections of art-historical imagery to demonstrate the formulaic recapitulation of motifs with significant gesture or pose,¹ as exemplified by Christ’s deposition from the cross in Figure 1. This is accompanied by a need for search engines that retrieve human figures with similar poses, facilitating the search for objects relevant to the individual scholar. It would thus become feasible to examine dominant pose types or time-dependent bodily phenomena on a large scale, as they were characteristic in Mannerism through the overlengthening of limbs, e.g., in Jacopo da Pontormo’s work (Figure 1b). Intra- as well as inter-iconographic recurrent motifs, whose radically altered semantics are disconcerting, might be thoroughly discussed in this context. To date, however, only few approaches exist for human pose estimation in art-historical images, possibly due to the lack of a sufficiently large domain-specific data set. To deal with this issue, one type of approaches uses pre-trained models, but without adapting them to the new domain [18, 29], while others apply style transfer to real-world data sets to obtain domain-specific training data [30], or fine-tune pre-trained models using small, keypoint-level annotated data sets [30].

In this paper, we propose a novel approach to quantitatively systematize the exploration of pose types in visual art utilizing semi-supervised learning. We suggest a two-stage approach based

¹For simplicity, we hereinafter do not distinguish between the terms ‘gesture,’ ‘posture,’ and ‘pose.’ Instead, we use the term ‘pose’ for any kind of physical expression.




Figure 1: The four depictions of Christ’s deposition from the cross highlight slightly varying poses: (a) Hans Pleydenwuff, 1465; (b) Pontorno, 1525–1528; (c) Caravaggio, 1603–1604; (d) Peter Paul Rubens, ca. 1612. All images are in the public domain.

on two Transformer models: the first model detects bounding boxes of human figures, while the second model predicts the keypoints of each box. We adapt a semi-supervised learning technique to reduce the performance loss caused by the shift between existing real-world data sets and the art domain, and to reduce the amount of art-historical labeled data. Our main contributions are as follows: (1) for object and keypoint detection, we suggest to combine semi-supervised pipelines through a two-step approach built on Transformer models with a teacher-student design; (2) to properly test our approach, we introduce a sufficiently large art-historical data set with both bounding box and keypoint annotations of human figures in 22 depiction styles; (3) in contrast to previous work, we show that the synthetic generation of seemingly ‘realistic’ art imagery inadequately reflects the stylistic diversity of historical artifacts. For both detection steps, the incorporation of manually labeled domain-specific material is performance-wise still required in the training and test phases. The code and models are available.²

The rest of the paper is structured as follows. Section 2 reviews related work on pose estimation and semi-supervised learning. In Section 3, we describe our pose estimator and its extension to a semi-supervised approach. In Section 4, we introduce our data sets and report on the ablation studies performed. Section 5 presents a user study to evaluate retrieval results from a human perspective. We conclude with Section 6 and outline areas of future work.

2 RELATED WORK

As with many other computer vision tasks, there has been steady progress in human pose estimation over recent years, particularly with the continued development of increasingly advanced deep learning models and self-supervised learning techniques.

Human pose estimation deals with the localization of a person’s skeleton by detecting associated keypoints, i.e., *skeleton coordinates* that mostly correspond to joint points of elbows, shoulders, etc. [7, 11, 24, 28, 41, 51, 55]. The problem can be solved in two ways. The *top-down* approach first detects persons, indexes them with bounding boxes, and then determines keypoints for each person [24, 28, 41]; while the *bottom-up* approach first detects keypoints, and then merges them to simultaneously identify persons and their basic pose [7, 11, 34]. Current work on the respective strategies shows that top-down methods generally yield better results, but at the cost of computational complexity [11]. Two-stage estimation makes the runtime linearly dependent on the number of detected persons in a scene, as the individual instances are cropped, and thus more forward steps are required for keypoint recognition.

²<https://github.com/TIBHannover/iart-semi-pose>, all last accessed on August 15, 2022.

However, since there is no real-time requirement for the domain considered here, runtime is of secondary importance. Further differences result from the prediction of the individual keypoints. Heatmap-based methods generate a dense likelihood map for the individual joints [41], whereas regression models directly predict coordinates of the individual components and optimize them [28]. While heatmap-based methods tend to perform better, the advantage of regression-based models is that they require fewer pre- and post-processing steps [28]. Therefore, we also make use of such models in our proposed method, as they are easier to integrate.

Few studies specifically address the estimation of human poses in art-historical images. This may be due to the fact that domain-specific data sets are usually only superficially indexed [2, 21, 33] and rarely include fine-grained annotations at the level of concrete image details [12, 31, 40, 54]. A publicly accessible data set that contains poses of human figures in artworks does not yet exist. Relevant previous work employs different approaches to deal with the lack of annotated training data: they (1) analyze only self-annotated data sets, without training models or performing inference [17]; (2) use trained pose estimators from another domain without adaptation [18, 29]; (3) apply style transfer to real-world data sets to close the domain gap [30]; or (4) leverage small, keypoint-level annotated data sets to fine-tune pre-trained models [30].

Semi-supervised learning aims to exploit a (potentially large) set of unlabeled data in addition to a (typically small) set of labeled data to improve the resulting model. To use the rest of the material during training, pseudo-labels are either generated [27, 52], or integrated into the loss with consistency regularization [26, 32]. One type of state-of-the-art methods uses a teacher-student approach. During training, an image is fed into a teacher model, which then generates a label for a student model that is being trained. The teacher model update can be iteratively selected from a previously trained student model [52], or the teacher is an Exponential Moving Average (EMA) of the student [43]. Another type of semi-supervised methods uses data augmentation to generate better feedback signals for unlabeled data, or combines pseudo-label generation and consistency regularization [3, 4, 39]. Similar to semi-supervised classification, localization methods are based on consistency regularization [19, 42] and pseudo-label generation [46, 53]. The challenge increases, however, since not only the respective concept must be assigned, but also its position in the image must be detected.

3 SEMI-SUPERVISED POSE ESTIMATION

In this section, we describe our method for automatic domain adaptation for human pose estimation. First, we introduce the two-stage

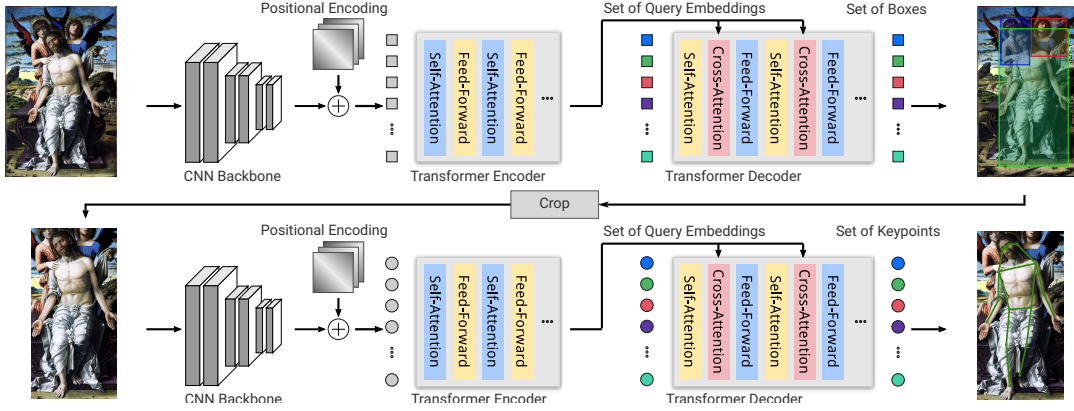


Figure 2: The two-stage human pose estimator uses two Transformers [8, 44] to predict human poses in an image. The input of the first model is the entire image, which, using a Convolutional Neural Network (CNN) backend and appropriate positional encoding, serves as input to a Transformer that predicts a fixed set of person bounding boxes. After filtering irrelevant detections, the individual boxes are cropped and serve as input for the second stage. This second Transformer model computes a set of keypoints that serve as the final prediction after filtering background classes.

Transformer-based detection model in Section 3.1. We then use it in the common approach of fine-tuning pre-trained models with stylized, approximately domain-specific images. In Section 3.2, we demonstrate how ‘real’ art-historical images can be used in the training stages with the extension of a semi-supervised process.

3.1 Transformer-based Detection

The proposed approach is organized in two steps: first, persons are detected in an input image and bounding boxes are computed; in a second step, the individual boxes are scanned for keypoints. The initial system is based on Li et al.’s method [28], which is built on two Transformer models for object detection [8, 44]. The overall architecture is shown in Figure 2.

In the **person detection phase**, feature descriptors are computed using a CNN backend combined with a two-dimensional position embedding. After this input is flattened into a sequence of visual features, it is passed to a Transformer encoder, which is later used in the cross-attention modules of the decoder. The other input of the Transformer decoder is a fixed set of trainable query embeddings, where the size of the set represents the maximum number of objects to be detected. The output is fed into two Multilayer Perceptron (MLP) heads. The first head acts as a classifier and distinguishes between person $c_{b,i}$ and background \emptyset , while the second one performs a regression on four outputs for the position and size of the corresponding box $b_i \in [0, 1]^4$. At the beginning of the **keypoint prediction stage**, visual features for each bounding box are determined using a CNN backend. The image features, combined with position encoding and a new set of input query embeddings, are transformed to a fixed set of keypoint predictions using the Transformer. The main difference between the two models is that the prediction head predicts only the coordinates of keypoints $k_i \in [0, 1]^2$, and instead of predicting only the person or background, classifies the type of keypoint $c_{k,i}$.

During the training phase, it is necessary to match the fixed set of predictions with the variable number of ground-truth labels

per image. We thus need to find an optimal assignment $\hat{\sigma}$ between prediction \hat{y} and ground-truth labels y in the permutation of N elements $\sigma \in \mathfrak{S}_N$ with the lowest matching cost L_m :

$$\hat{\sigma} = \arg \min_{\sigma \in \mathfrak{S}_N} \sum_i^N L_m(y_i, \hat{y}_{\sigma(i)}) \quad (1)$$

The optimal solution for this problem can be solved using the Hungarian algorithm [25] and yields the assignment function $\hat{\sigma}(i)$. The assignment loss includes both the class probability and the position of the predicted object compared to the ground-truth annotation. For bounding box prediction with index $\sigma(i)$, we define the class probability $c_{b,i}$ as $\hat{p}_{\sigma(i)}(c_{b,i})$ and the predicted box as $\hat{b}_{\sigma(i)}$. Similarly, for keypoint prediction, we define the probability of class $c_{k,i}$ as $\hat{p}_{\sigma(c_{k,i})}(i)$ and the predicted keypoint as $\hat{k}_{\sigma(i)}$. With these definitions, we establish the following loss functions:

$$L_{m,b}(y, \hat{y}) = -\mathbb{1}_{\{c_{b,i} \neq \emptyset\}} \hat{p}_{\sigma(i)}(c_{b,i}) + \mathbb{1}_{\{c_{b,i} \neq \emptyset\}} L_b(b_i, \hat{b}_{\sigma(i)}) \quad (2)$$

$$L_{m,k}(y, \hat{y}) = -\mathbb{1}_{\{c_{k,i} \neq \emptyset\}} \hat{p}_{\sigma(i)}(c_{k,i}) + \mathbb{1}_{\{c_{k,i} \neq \emptyset\}} L_k(k_i, \hat{k}_{\sigma(i)}) \quad (3)$$

For bounding box prediction, the class probability defined as the $L1$ -distance of the bounding box b_i , and the Generalized Intersection over Union (GIoU) [37] $L_{iou}(\cdot, \cdot)$ are chosen as the basis for cost function L_b , where we follow Li et al.’s approach and implementation [28]. For keypoints k_i , only the class probability and the $L1$ -distance of the coordinates are considered:

$$L_b(b_i, \hat{b}_{\sigma(i)}) = \lambda_{iou} L_{iou}(b_i, \hat{b}_{\sigma(i)}) + \lambda_{L1} \|b_i - \hat{b}_{\sigma(i)}\| \quad (4)$$

$$L_k(k_i, \hat{k}_{\sigma(i)}) = \lambda_{L1} \|k_i - \hat{k}_{\sigma(i)}\| \quad (5)$$

where hyperparameters λ_{iou} and λ_{L1} indicate the weight of each loss component. Predictions that could not be assigned to a ground-truth label are instead assigned to the background class \emptyset during optimization; their bounding boxes and keypoint coordinates are not considered in the loss. After the best assignment is found, the

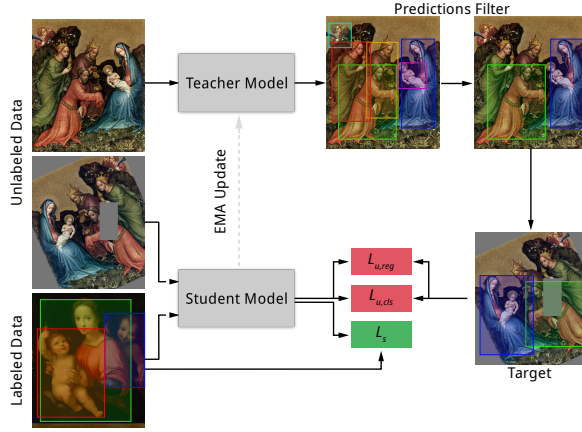


Figure 3: In the semi-supervised training pipeline adapted from Xu et al. [53], each batch consists of labeled and unlabeled images with strong and weak augmentations generated for unlabeled ones. The teacher uses the weakly labeled data to generate pseudo bounding boxes (or pseudo keypoints) that are used to train the strongly augmented images. This involves thresholding the predictions and then transferring the corresponding boxes (or keypoints) to the coordinate system of the strongly augmented image.

loss can be calculated as follows:

$$L_{H,b}(y, \hat{y}) = \sum_{i=1}^N \left[-\log \hat{p}_{\hat{\sigma}}(c_{b,i}) + \mathbb{1}_{\{c_{b,i} \neq 0\}} L_b(b_i, \hat{b}_{\hat{\sigma}}(i)) \right] \quad (6)$$

$$L_{H,k}(y, \hat{y}) = \sum_{i=1}^N \left[-\log \hat{p}_{\hat{\sigma}}(c_{k,i}) + \mathbb{1}_{\{c_{k,i} \neq 0\}} L_k(b_i, \hat{b}_{\hat{\sigma}}(i)) \right] \quad (7)$$

During the inference of bounding box prediction, it is sufficient to filter the predicted boxes using a threshold function. However, during the inference step of keypoint prediction, it is necessary to find an optimal assignment again because the Transformer model predicts up to N points, but the number is usually larger than the maximum number of possible keypoints per person. Since no ground-truth information is known during inference, the following cost function is used:

$$L_{m,k}(y, \hat{y}) = -\hat{p}_{\sigma(i)}(c_{k,i}) \quad (8)$$

Compared to object detection methods such as Faster Region-based CNN (R-CNN) [36] and YOLO (You Only Look Once) [35], the approach does not predict multiple bounding box candidates for each image region, but only a fixed set of boxes for each image. This greatly simplifies post-processing, as no overlapping bounding boxes are predicted for same-person instances, and the imbalance between background and foreground classes is much smaller.

3.2 Semi-supervised Domain Adaptation

To extend the available data sets for bounding box and keypoint detection in art-historical images, we augment the training pipeline by adapting the semi-supervised approach from Xu et al. [53]. Since we use a Transformer model instead of a Faster R-CNN, the number

of predicted bounding boxes and keypoints is considerably smaller and simplifies certain steps. An overview of the semi-supervised pipeline is shown in Figure 3. The basic principle is to use both labeled and unlabeled examples to train a student model. Here, the teacher, whose weights are based on the EMA of the student weights, serves as a generator of pseudo-labels for bounding boxes and keypoints. For this purpose, weakly augmented unlabeled images are used for person detection and weakly augmented cropped images are used for keypoint prediction. Subsequently, the predicted objects are filtered with the threshold $\tau = 0.9$ and projected onto the strongly augmented unlabeled images. Contrary to Xu et al. [53], it is not possible to determine target labels for the background class from the teacher, because negative teacher predictions do not have to contain any valid coordinates and therefore cannot be assigned to an output of the student using the Hungarian algorithm. Therefore, we use the teacher prediction for bounding boxes and keypoints only if it is not a background class. To not distort the ratio between negative and positive boxes or keypoints, we suggest to use the same threshold to filter negative examples; but this time from the forward step of the student. This is necessary because there is no relationship between the predicted coordinates of the teacher’s negative classes and the student’s negative predictions. The total loss now includes a supervised component L_s and an unsupervised component L_u . It is calculated as follows:

$$L = L_s + \lambda_u L_u \quad (9)$$

Depending on the current target, the supervised loss is the same as for supervised learning, $L_s \in \{L_{H,b}, L_{H,k}\}$. For the unsupervised loss part, we use the prediction of the teacher model to detect bounding boxes or keypoints. Therefore, for the prediction of the bounding box with index i , we define the probability of class $c_{b,i}$ as $\hat{p}^t(c_{b,i})$ and the predicted box as \hat{b}^t . Similarly, for the teacher keypoint prediction, we define the probability of class $c_{k,i}$ as $\hat{p}^t(c_{k,i})$ and the predicted keypoint as \hat{k}^t . With these definitions, we can establish the loss functions:

$$L_{u,reg,b} = \sum_i^N \mathbb{1}_{\{c_{b,i} \neq 0; \hat{p}^t(c_{b,i}) \geq \tau\}} L_b(\hat{b}_i^t, \hat{b}_{\hat{\sigma}}(i)) \quad (10)$$

$$L_{u,reg,k} = \sum_i^N \mathbb{1}_{\{c_{k,i} \neq 0; \hat{p}^t(c_{k,i}) \geq \tau\}} L_k(\hat{k}_i^t, \hat{k}_{\hat{\sigma}}(i)) \quad (11)$$

The classification loss of the unlabeled examples is given by the positive classes resulting from the teacher’s probability of exceeding threshold τ and the negative examples from the student’s prediction:

$$L_{u,cls,b} = - \sum_i^N \mathbb{1}_{\{c_{b,i} \neq 0; \hat{p}^t(c_{b,i}) \geq \tau\}} \log \hat{p}_{\hat{\sigma}}(c_{b,i}) \quad (12)$$

$$- \sum_i^N \mathbb{1}_{\{c_{b,i} = 0; \hat{p}_{\hat{\sigma}}(c_{b,i}) \geq \tau\}} \log \hat{p}_{\hat{\sigma}}(c_{b,i})$$

$$L_{u,cls,k} = - \sum_i^N \mathbb{1}_{\{c_{k,i} \neq 0; \hat{p}^t(c_{k,i}) \geq \tau\}} \log \hat{p}_{\hat{\sigma}}(c_{k,i}) \quad (13)$$

$$- \sum_i^N \mathbb{1}_{\{c_{k,i} = 0; \hat{p}_{\hat{\sigma}}(c_{k,i}) \geq \tau\}} \log \hat{p}_{\hat{\sigma}}(c_{k,i})$$

Table 1: An overview is given of the data sets used in our experiments. Persons are indicated by bounding boxes associated with them. Up to 17 keypoints are stored per person.

Data set	Split	Images	Persons	Keypoints
COCO 2017	Training	118,287	262,465	1,642,283
	Validation	5,000	11,004	68,215
	Test	0	0	0
	Total	123,287	273,469	1,710,498
COCO 2017 (stylized)	Training	236,574	524,930	3,284,566
	Validation	10,000	22,008	136,430
	Test	0	0	0
	Total	246,574	546,938	3,420,996
People-Art	Training	1,746	1,330	0
	Validation	1,489	1,080	0
	Test	1,616	1,088	0
	Total	4,851	3,498	0
PoPArt	Training	1,553	2,069	30,415
	Validation	643	704	10,367
	Test	663	741	10,863
	Total	2,859	3,514	51,645
ART500k	Training	318,869	0	0
	Validation	0	0	0
	Test	0	0	0
	Total	318,869	0	0

4 EXPERIMENTAL SETUP AND RESULTS

In this section, we introduce our data sets and discuss the quantitative and qualitative studies. For the training and test phases of our pipelines, we use various real-world, synthetically generated, and art-historical data sets (Section 4.1). To evaluate the performance of each model and approach, we first conduct a series of ablation studies (Section 4.2) and then qualitatively assess our method’s ability to provide reasonable predictions (Section 4.3). To evaluate the experiments, we use the metrics and tools from the COCO API.³

4.1 Data Sets

An overview of the data sets used in our experiments with their respective splits is shown in Table 1. All data sets are based on the Common Objects in Context (COCO) format, where each person instance is labeled with up to 17 keypoints.

The largest annotated data set results from the **COCO 2017** detection and keypoint challenge, which includes 118,287 training and 5,000 validation images with person instances.⁴ To evaluate the performance of the common scenario that uses style transfer to close the domain gap between annotated real-world training and art-historical inference data, we generate a stylized version of the data set. For this purpose, we leverage the style transfer approach

from Chen et al. [10] to create two style variants for each COCO image, where the style images are randomly selected from the Painter by Numbers data set [33].

The models are grounded in two domain-specific, sufficiently large data sets that recycle openly licensed subsets of the art-historical online encyclopedia WikiArt⁵: the 2016 compiled People-Art data set [6, 50], in which human figures are marked with bounding boxes enclosing them. The second data set, called Poses of People in Art (PoPArt), is introduced here and identifies 17 limb points in addition to bounding boxes. Both data sets approximately reflect the diversity of art-historical depictions of human figures through time by featuring 43 and 22 different styles, respectively; ranging from impressionistic to neo-figurative and realistic variants. The pre-existing **People-Art** data set is enhanced on two levels. First, we integrate additional negative examples of mammals that were frequently false positively classified as humans [50]. Second, we use the largest resolution of images provided by WikiArt to avoid further complicating the detection of relatively small figures due to possible image artifacts in low-resolution reproductions. After these preparatory measures, People-Art features 1,746 training, 1,489 validation, and 1,616 test images. The annotation of the novel **PoPArt** data set was performed according to the following principles (see Figure 5a for some examples with ground-truth annotations): (1) the body of a human figure must be recognizable, which implies that more than six keypoints are annotatable, covering at least head and shoulder area; (2) a maximum of four figures are annotated per image; if more than four instances are shown, those whose body permits to annotate as many limbs as possible are selected; (3) if an occluded body part can be sufficiently approximated by another visible one, the respective associated keypoint is annotated; (4) in profile views, eyes and ears are usually annotated on the non-visible side of the face as well. The data set includes 1,553 training, 643 validation, and 663 test images, where each split contains proportionally the same number of images per style.

With the **ART500k** data set [31], we moreover integrate an art-historical data set not annotated with person instances into the training procedure. A 50% split of all ART500k images with a total of 318,869 examples is generated, which we use in our semi-supervised learning approach as unlabeled data.

4.2 Ablation Study

For **person detection**, we leverage the weights of a Detection Transformer (DETR) model [8] pre-trained on COCO 2017 and reinitialize the classification head. An Adam optimizer [22] with a learning rate of $lr = 5e - 6$ is used for the Transformer and with $lr = 1e - 7$ for the ResNet-50 backbone [15]. Similar to Li et al. [28], all classes except persons are ignored; small bounding boxes are not considered. Models are trained for 200,000 iterations with a batch size of four, with all images randomly scaled to a maximum size of 1,333 pixels per side. When training the semi-supervised models, the batch size is increased by four additional unlabeled images. The weights of the different loss hyperparameters are set to $\lambda_{L1} = 5$, $\lambda_{iou} = 2$, and $\lambda_u = 0.5$.

The results for the respective test sets are shown in Table 2, including a comparison with the best state-of-the-art method to

³<https://github.com/cocodataset/cocoapi>.

⁴<https://www.kaggle.com/datasets/awsaf49/coco-2017-dataset>.

⁵<https://www.wikiart.org/>.

Table 2: Person detection results are reported for the People-Art and PoPart test sets, respectively. For PoPart, AP_S is neglected as no test data is available for small human figures, most of which have no annotatable pose due to their size. Entries without style transfer and without semi-supervised learning correspond to the state-of-the-art method of Li et al. [28] with fine-tuning to the respective training data set. The best performing approach per test set is indicated in bold.

Test set	Train set	Stylized	Semi	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L	AR
People-Art	COCO 2017	0 %		0.3118	0.5106	0.3175	0.0075	0.2118	0.3294	0.6728
	COCO 2017	0 %	✓	0.3696	0.5970	0.3885	0.0007	0.2115	0.3950	0.7351
	COCO 2017	50 %		0.3686	0.6113	0.3871	0.0045	0.2386	0.3941	0.7257
	COCO 2017	50 %	✓	0.3744	0.6277	0.3792	0.0024	0.2193	0.4011	0.7296
	COCO 2017	100 %		0.3727	0.6256	0.3922	0.0240	0.2406	0.3981	0.7165
	COCO 2017	100 %	✓	0.3846	0.6333	0.4047	0.0115	0.2313	0.4108	0.7221
	People-Art	0 %		0.4280	0.7279	0.4350	0.0676	0.2123	0.4636	0.7041
	People-Art	0 %	✓	0.4428	0.7381	0.4590	0.0509	0.2412	0.4769	0.7291
PoPart	COCO 2017	0 %		0.2287	0.3041	0.2433		0.1096	0.2336	0.7997
	COCO 2017	0 %	✓	0.2422	0.3353	0.2612		0.0324	0.2469	0.8377
	COCO 2017	50 %		0.2322	0.3168	0.2480		0.0400	0.2397	0.8365
	COCO 2017	50 %	✓	0.2261	0.3125	0.2452		0.0347	0.2324	0.8277
	COCO 2017	100 %		0.2542	0.3540	0.2730		0.0360	0.2624	0.8128
	COCO 2017	100 %	✓	0.2359	0.3310	0.2516		0.0480	0.2423	0.8284
	PoPart	0 %		0.4898	0.6566	0.5279		0.2639	0.4945	0.8468
	PoPart	0 %	✓	0.5073	0.6728	0.5302		0.2132	0.5119	0.8561

date [28]. We notice that our semi-supervised learning technique on People-Art always results in an improvement of Average Precision (AP) and Average Recall (AR). Moreover, AP maintains this advantage as the proportion of style-transferred material increases, but becomes successively smaller. The domain-specific data further increases the performance significantly, such that AP rises from 0.4280 to 0.4428 and AR from 0.7041 to 0.7291. With $AP_{50} = 0.7381$, the performance of our approach is considerably above the best results of $AP_{50} = 0.68$ and $AP_{50} = 0.583$ reported so far by Kadish et al. [20] and Gonthier et al. [13] for the data set, respectively. For PoPart, we find that semi-supervised learning with art-historical images enhances AP less; thus, our proposed method with COCO 2017 annotations has similar performance to using style transfer. The comparison between training with COCO 2017 data and training on PoPart indicates a larger improvement especially in AP. This deviation can be explained by the different types of annotations, as PoPart was annotated exclusively for pose estimation and contains fundamentally fewer ground-truth bounding boxes of human figures. Nevertheless, our proposed semi-supervised learning approach is beneficial: the performance increases from 0.4898 to 0.5073 for AP and from 0.8468 to 0.8561 for AR.

In the **keypoint prediction** stage, we use the High-Resolution Net with 32 feature channels (HRNet-W32) as backbone with an input resolution of 384×288 pixels [41]. Again, we leverage the pre-trained weights on COCO 2017 from Li et al. [28] and reinitialize the classification layer. The model is trained for 150,000 iterations with a batch size of 16; the learning rates are set to $lr = 1e - 5$ for the Transformer and $lr = 1e - 6$ for the HRNet. We divide both rates by 10 and train for another 50,000 iterations with Adam. For the semi-supervised methods, we add to the batch 16 unlabeled images generated from the models' predictions from Table 2 on ART500k.

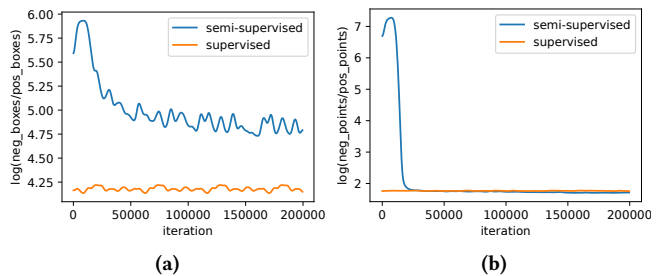


Figure 4: The distribution of positive and negative classes on PoPart (orange) and the teacher's predicted distribution for unlabeled data on ART500k (blue) are shown. It is evident that the teacher recognizes fewer bounding boxes in the person detection phase (a) and estimates more points in the keypoint prediction phase (b) in comparison.

Predicted bounding boxes whose confidence level is above 0.5 are used for this purpose. The effects of keypoint prediction are similar to those of person detection: we observe that AR can be significantly improved by our semi-supervised learning technique. Models not only trained with style-transferred images show an increase in AP. In particular, for those using PoPart, AP rises from 0.4844 to 0.5258 and AR from 0.7078 to 0.7464. Results for the PoPart test set are summarized in Table 3. Unlike Jeníček and Chum [18], we find that OpenPose [7] with $AP = 0.1388$ and $AR = 0.4382$ is not competitive to approaches that are specifically trained for the given task.

To evaluate the behavior of our **semi-supervised approach**, we examine the number of positive and negative teacher predictions during training. To this end, we illustrate the ratios between

Table 3: Keypoint detection results are reported for the PoPArt test set with predicted bounding boxes of the model with the same strategy. AP_S is neglected as no test data is available for small human figures, most of which have no annotatable pose due to their size. For PoPArt train sets, the first entry refers to the training data set used for bounding box detection and the second to the training data set used for keypoint prediction. The best performing approach is indicated in bold.

Test set	Train set	Stylized	Semi	AP	AP_{50}	AP_{75}	AP_M	AP_L	AR
PoPArt	COCO 2017	0 %		0.2285	0.2811	0.2545	0.0236	0.2367	0.5540
	COCO 2017	0 %	✓	0.2525	0.3173	0.2810	0.0122	0.2639	0.7009
	COCO 2017	50 %		0.2401	0.3072	0.2672	0.0215	0.2531	0.6672
	COCO 2017	50 %	✓	0.2413	0.3052	0.2665	0.0180	0.2554	0.6880
	COCO 2017	100 %		0.2657	0.3426	0.2932	0.0153	0.2845	0.6765
	COCO 2017	100 %	✓	0.2518	0.3167	0.2813	0.0169	0.2653	0.6896
	People-Art/PoPArt	0 %		0.2841	0.3622	0.3073	0.0378	0.2916	0.7185
	People-Art/PoPArt	0 %		0.2971	0.3637	0.3272	0.0204	0.3118	0.7583
	PoPArt/PoPArt	0 %		0.4844	0.6060	0.5319	0.0771	0.4920	0.7078
	PoPArt/PoPArt	0 %		0.5258	0.6392	0.5735	0.0308	0.5350	0.7464

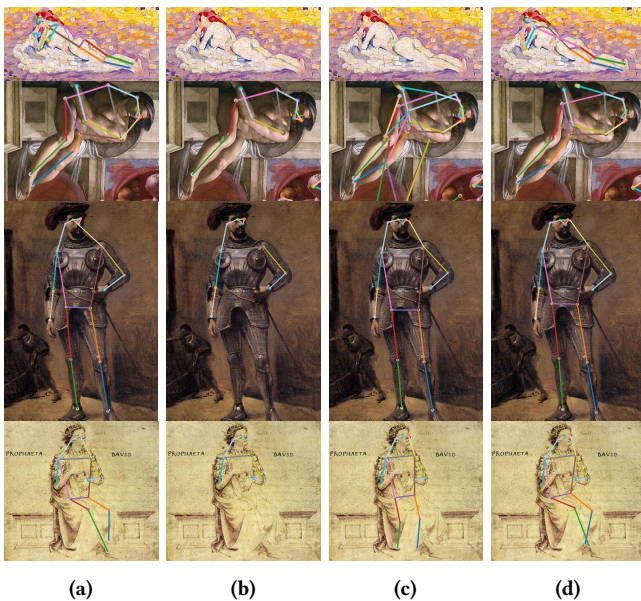


Figure 5: Ground-truth annotations (a), as well as predictions of OpenPose [7] (b), the model trained without style transfer and without semi-supervised learning (c), and PoPArt (d) overlaid on test examples from PoPArt.

negative and positive bounding boxes (Figure 4a) and keypoints (Figure 4b) of the labeled and unlabeled parts of a batch. As we compute the target labels for the background class directly from the student’s predictions, we can see in both cases that the ratio of the background increases sharply until it reaches the maximum at about 10,000 iterations. After that, it starts to decrease in favor of positive classes as the confidence score of the teacher’s predictions starts to exceed threshold τ . In case of keypoints, it becomes apparent that the ratio between supervised (PoPArt) and unsupervised components (ART500k) per batch is equalized, and later on average more keypoints are detected in ART500k than in PoPArt.

4.3 Qualitative Analysis

To qualitatively assess our method’s ability to provide reasonable predictions, we visually compare it to ground-truth annotations and two of the other models. Figure 5b illustrates that OpenPose almost consistently tends to estimate only parts of the face and some points of the torso; holistically correct predictions are rare. Bodies in non-realistic settings are often not captured, exemplified by Henri Edmond Cross’s Neo-Impressionist example from the early 20th century (Figure 5b, *first row*). This is also noticeable in the model trained on COCO 2017 without style transfer and unsupervised learning (Figure 5c). However, more suitable approximations of the lower body are identified, at least for Jean-Baptiste Camille Corot’s *Knight* (1868; Figure 5c, *third row*) and Fra Angelico’s religious drawing of King David (ca. 1430; *fourth row*). Highly problematic, though, are hidden limbs or bodies not depicted from usual perspectives, illustrated by the detail of Michelangelo’s Sistine Chapel ceiling painting in Figure 5c (*second row*). Our proposed model, trained on PoPArt and with semi-supervised learning, even manages predominantly complex scenarios (Figure 5d). Minor errors result from limbs assigned to the wrong side of the body (*first row*), poorly contrasting or rather abstractly drawn body parts—or overlaps with limbs of other persons, which in PoPArt were primarily due to Aubrey Beardsley’s works. This is especially true for styles that introduce complications even when manually labeled, e.g., in case of the Japanese genre Ukiyo-e, since expressive poses with strongly flowing robes often lack clear assignment of joint points. In addition, the correct assignment of points can be disturbed if the image shows a person and his or her mirror image.

5 USER STUDY ON RETRIEVAL RESULTS

In this section, we report the results of a user study that aimed to evaluate the quality of the automatically generated keypoints from a human perspective in a retrieval scenario. We first describe the generation of keypoint descriptors and the experimental setup of the user study before discussing the results.

Keypoint descriptors. For the retrieval task, we convert keypoints into a consistent feature vector representation. In doing so,

Table 4: Results of the user study are reported on the retrieval of similar poses with Normalized Discounted Cumulative Gain (NDCG) as the ranking metric.

Train set	Stylized	Semi	@5	@10	@15
COCO 2017	0 %		0.5626	0.5929	0.6309
COCO 2017	0 %	✓	0.5702	0.5676	0.5957
COCO 2017	50 %		0.6124	0.6054	0.6234
COCO 2017	50 %	✓	0.5713	0.5900	0.6094
COCO 2017	100 %		0.5728	0.5958	0.6131
COCO 2017	100 %	✓	0.5845	0.6069	0.6304
PoPArT	0 %		0.5675	0.5722	0.5943
PoPArT	0 %	✓	0.6413	0.6205	0.6344



Figure 6: Query images for the user study include art-historical poses such as ‘Adlocutio’ and ‘Venus pudica.’

descriptors for the same pose should be nearly identical regardless of position or scale. As pose discrimination depends heavily on the relational configuration between body parts [16], we do not leverage joint coordinates directly [14, 38]. Instead, we build on Chen et al. [9] and employ a 52-dimensional feature descriptor that uses the orientation between two keypoints. We obtain 1,515 images from the ART500k data set not used for training in Section 4.2, to which bounding box and keypoint models are applied. For each pose, the descriptor from Chen et al. [9] is calculated. In addition, we selected 10 poses with varying art-historical specificity and utilized them as query images (Figure 6). The small number of examples naturally can only inadequately cover the large variability of relevant body constellations; it is, nonetheless, sufficient to ascertain the models’ basic suitability for retrieval tasks.

Experimental setup. For our study, we developed a web interface with detailed instructions for annotation. A total of 12 subjects were recruited, personally invited by the participating departments of computer science and art history. These included seven computer scientists, two art historians, and three persons from other professions. In the study, several pages were shown, consisting of a query image and the corresponding top-20 retrieval results. For each displayed image, participants were asked to vote on whether they thought it was ‘relevant,’ ‘irrelevant,’ or ‘indifferent’ to the query. After the questioning, the individual results were ranked in this order: ‘relevant,’ ‘indifferent,’ and ‘irrelevant.’ We used Euclidean distance to compute a ranking based on the automatically

computed descriptors and compared it to the user-generated ranking. The results of the user study are reported in Table 4 and show that our proposed approach also outperforms competing models in retrieval, nevertheless, with decreasing variations between models. This can possibly be explained by the fact that it is not necessarily relevant for a user if the alignment of individual keypoints changes as long as the basic pose has very similar meaning. However, it may also be that the number of subjects is too small for such conclusions, or that the participants’ art-historical knowledge was insufficient to interpret certain details of the poses. In this context, the degree of similarity at which subjects consider poses to be similar is relevant. For instance, one participant excluded crucifixion scenes in which Christ looked to the left rather than downward with his head bowed, as in the query image.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we have investigated domain adaption techniques to estimate human poses in art-historical images. To this end, we have suggested a two-stage approach based on two Transformer models that utilizes a semi-supervised teacher-student design. To reduce the gap between photographs of real-world objects and the art domain, we augmented images depicting real-world scenes with unlabeled, domain-specific data. Moreover, we introduced a reasonably large art-historical data set called Poses of People in Art (PoPArT) to systematically test the validity of human pose estimators. Comparisons with more common approaches that use pre-trained models or adapt existing data sets with style transfer indicated that performance can be further improved with unlabeled data. While it is not necessary to annotate large amounts of art-historical data, it is essential to include at least smaller, domain-specific labeled data in the training procedure, rather than relying solely on synthetically generated imagery. Depending on the test set, models trained entirely or partially with style transfer underperform in Average Precision by between 7.32 to 28.12 % for person detection and between 27.33 to 28.15 % for keypoint prediction, even with semi-supervised learning. Furthermore, a user study confirmed the feasibility of the proposed approach for retrieval tasks, thus also enabling the search for resembling poses of human figures; however, in this case the difference with other models performance-wise is smaller. Our method enables the engagement of humanities scholars by providing them with state-of-the-art methods for indexing human poses in large art image databases. Although our approach specifically targets the curation of art and cultural objects, it is likely applicable to other domains with few labeled training data.

In the future, we intend to analyze the potential of recently introduced Transformer models, such as the Pyramid Vision Transformer presented by Wang et al. [47]. Further improvement of the training process could be achieved by applying style transfer to unlabeled instead of only labeled data. We also plan to extend the PoPArT data set with additional bounding boxes, enhancing its usefulness for training person detection models.

ACKNOWLEDGEMENTS

This work was partly funded by the German Research Foundation (DFG) under project number 415796915. We thank the participants in the retrieval study for their valuable contributions.

REFERENCES

- [1] Moshe Barasch. 1987. *Giotto and the Language of Gesture*. Cambridge University Press.
- [2] Matthias Becker, Martin Bogner, Fabian Bross, François Bry, Caterina Campanella, Laura Commare, Silvia Cramerotti, Katharina Jakob, Martin Josko, Fabian Kneißl, Hubertus Kohle, Thomas Krefeld, Elena Levushkina, Stephan Lücke, Alessandra Puglisi, Anke Regner, Christian Riepl, Clemens Schefels, Corina Schemainda, Eva Schmidt, Stefanie Schneider, Gerhard Schön, Klaus Schulz, Franz Sigmüller, Bartholomäus Steinmayr, Florian Störkle, Iris Teske, and Christoph Wieser. 2018. *ARTigo – Social Image Tagging [Dataset and Images]*. <https://doi.org/10.5282/ubm/data.136>
- [3] David Berthelot, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. 2020. ReMixMatch: Semi-Supervised Learning with Distribution Matching and Augmentation Anchoring. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020*.
- [4] David Berthelot, Nicholas Carlini, Ian J. Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. 2019. MixMatch: A Holistic Approach to Semi-Supervised Learning. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.). 5050–5060.
- [5] Richard Brilliant. 1963. *Gesture and Rank in Roman Art. The Use of Gestures to Denote Status in Roman Sculpture and Coinage*. Connecticut Academy of Arts & Sciences.
- [6] Hongping Cai, Qi Wu, Tadeo Corradi, and Peter Hall. 2015. The Cross-Depiction Problem: Computer Vision Algorithms for Recognising Objects in Artwork and in Photographs. (2015). [arXiv:1505.00110](https://arxiv.org/abs/1505.00110)
- [7] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2021. OpenPose: Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 1 (2021), 172–186. <https://doi.org/10.1109/TPAMI.2019.2929257>
- [8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-End Object Detection with Transformers. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12346)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 213–229. https://doi.org/10.1007/978-3-030-58452-8_13
- [9] Cheng Chen, Yueting Zhuang, Feiping Nie, Yi Yang, Fei Wu, and Jun Xiao. 2011. Learning a 3D Human Pose Distance Metric from Geometric Pose Descriptor. *IEEE Trans. Vis. Comput. Graph.* 17, 11 (2011), 1676–1689. <https://doi.org/10.1109/TVCG.2010.272>
- [10] Haibo Chen, Lei Zhao, Zhizhong Wang, Zhang Hui Ming, Zhiwen Zuo, Ailin Li, Wei Xing, and Dongming Lu. 2021. Artistic Style Transfer with Internal-external Learning and Contrastive Learning. In *Thirty-Fifth Conference on Neural Information Processing Systems*.
- [11] Bowen Cheng, Bin Xiao, Jingdong Wang, Honghui Shi, Thomas S. Huang, and Lei Zhang. 2020. HigherHRNet: Scale-Aware Representation Learning for Bottom-Up Human Pose Estimation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020*. Computer Vision Foundation / IEEE, 5385–5394. <https://doi.org/10.1109/CVPR42600.2020.00543>
- [12] Noa Garcia and George Vogiatis. 2018. How to Read Paintings: Semantic Art Understanding with Multi-modal Retrieval. In *Computer Vision - ECCV 2018 Workshops - Munich, Germany, September 8–14, 2018, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 11130)*, Laura Leal-Taixé and Stefan Roth (Eds.). Springer, 676–691. https://doi.org/10.1007/978-3-030-11012-3_52
- [13] Nicolas Gonthier, Saïd Ladjal, and Yann Gousseau. 2022. Multiple Instance Learning on Deep Features for Weakly Supervised Object Detection with Extreme Domain Shifts. *Comput. Vis. Image Underst.* 214 (2022), 103299. <https://doi.org/10.1016/j.cviu.2021.103299>
- [14] Tatsuya Harada, So Taoka, Taketoshi Mori, and Tomomasa Sato. 2004. Quantitative Evaluation Method for Pose and Motion Similarity Based on Human Perception. In *4th IEEE/RAS International Conference on Humanoid Robots, Humanoids 2004, Santa Monica, CA, USA, November 10–12, 2004*. IEEE, 494–512. <https://doi.org/10.1109/ICHR.2004.1442140>
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27–30, 2016*. IEEE Computer Society, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [16] Edmond S. L. Ho and Taku Komura. 2009. Indexing and Retrieving Motions of Characters in Close Contact. *IEEE Trans. Comput. Graph.* 15, 3 (2009), 481–492. <https://doi.org/10.1109/TVCG.2008.199>
- [17] Leonardo Impett and Sabine Süsstrunk. 2016. Pose and Pathosformel in Aby Warburg's Bilderatlas. In *Computer Vision - ECCV 2016 Workshops - Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 9913)*, Gang Hua and Hervé Jégou (Eds.). Springer, 888–902. https://doi.org/10.1007/978-3-319-46604-0_61
- [18] Tomás Jeníček and Ondrej Chum. 2019. Linking Art through Human Poses. In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019, Sydney, Australia, September 20–25, 2019*. IEEE, 1338–1345. <https://doi.org/10.1109/ICDAR.2019.00216>
- [19] Jisoo Jeong, Seungeui Lee, Jeeseo Kim, and Nojun Kwak. 2019. Consistency-based Semi-supervised Learning for Object detection. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.). 10758–10767.
- [20] David Kadish, Sebastian Risi, and Anders Sundnes Løvlie. 2021. Improving Object Detection in Art Images Using Only Style Transfer. In *International Joint Conference on Neural Networks, IJCNN 2021, Shenzhen, China, July 18–22, 2021*. IEEE, 1–8. <https://doi.org/10.1109/IJCNN52387.2021.9534264>
- [21] Sergey Karayev, Matthew Trentacoste, Helen Han, Aseem Agarwala, Trevor Darrell, Aaron Hertzmann, and Holger Winnemoeller. 2014. Recognizing Image Style. In *British Machine Vision Conference, BMVC 2014, Nottingham, UK, September 1–5, 2014*, Michel François Valstar, Andrew P. French, and Tony P. Pridmore (Eds.). BMVA Press.
- [22] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6980>
- [23] James R. Knowlson. 1965. The Idea of Gesture as a Universal Language in the XVIIIth and XVIIIth Centuries. *Journal of the History of Ideas* 26 (1965), 495–508.
- [24] Sven Kreiss, Lorenzo Bertoni, and Alexandre Alahi. 2019. PiPaF: Composite Fields for Human Pose Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019*. Computer Vision Foundation / IEEE, 11977–11986. <https://doi.org/10.1109/CVPR.2019.01225>
- [25] Harold W. Kuhn. 1955. The Hungarian Method for the Assignment Problem. *Naval research logistics quarterly* 2, 1–2 (1955), 83–97.
- [26] Samuli Laine and Timo Aila. 2017. Temporal Ensembling for Semi-Supervised Learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*.
- [27] Dong-Hyun Lee. 2013. Pseudo-label: The Simple and Efficient Semi-supervised Learning Method for Deep Neural Networks. In *Workshop on Challenges in Representation Learning, ICML, Vol. 3*. 896.
- [28] Ke Li, Shijie Wang, Xiang Zhang, Yifan Xu, Weijian Xu, and Zhuowen Tu. 2021. Pose Recognition With Cascade Transformers. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19–25, 2021*. Computer Vision Foundation / IEEE, 1944–1953.
- [29] Prathmesh Madhu, Tilman Marquart, Ronak Kosti, Peter Bell, Andreas K. Maier, and Vincent Christlein. 2020. Understanding Compositional Structures in Art Historical Images Using Pose and Gaze Priors - Towards Scene Understanding in Digital Art History. In *Computer Vision - ECCV 2020 Workshops - Glasgow, UK, August 23–28, 2020, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 12536)*, Adrien Bartoli and Andrea Fusiello (Eds.). Springer, 109–125. https://doi.org/10.1007/978-3-030-66096-3_9
- [30] Prathmesh Madhu, Angel Villar-Corralles, Ronak Kosti, Torsten Bendschus, Corinna Reinhardt, Peter Bell, Andreas K. Maier, and Vincent Christlein. 2020. Enhancing Human Pose Estimation in Ancient Vase Paintings via Perceptually-grounded Style Transfer Learning. (2020). [arXiv:2012.05616](https://arxiv.org/abs/2012.05616)
- [31] Hui Mao, Ming Cheung, and James She. 2017. DeepArt: Learning Joint Representations of Visual Arts. In *Proceedings of the 2017 ACM on Multimedia Conference, MM 2017, Mountain View, CA, USA, October 23–27, 2017*, Qiong Liu, Rainer Lienhart, Haohong Wang, Sheng-Wei "Kuan-Ta" Chen, Susanne Boll, Yi-Ping Phoebe Chen, Gerald Friedland, Jia Li, and Shuicheng Yan (Eds.). ACM, 1183–1191. <https://doi.org/10.1145/3123266.3123405>
- [32] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2019. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning. *IEEE Trans. Pattern Anal. Mach. Intell.* 41, 8 (2019), 1979–1993. <https://doi.org/10.1109/TPAMI.2018.2858821>
- [33] Kiri Nichol. 2016. *Painter by Numbers*. Retrieved June 15, 2021 from <https://www.kaggle.com/c/painter-by-numbers>
- [34] George Papandreou, Tyler Zhu, Liang-Chieh Chen, Spyros Gidaris, Jonathan Tompson, and Kevin Murphy. 2018. PersonLab: Person Pose Estimation and Instance Segmentation with a Bottom-Up, Part-Based, Geometric Embedding Model. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8–14, 2018, Proceedings, Part XIV (Lecture Notes in Computer Science, Vol. 11218)*, Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (Eds.). Springer, 282–299. https://doi.org/10.1007/978-3-030-01264-9_17
- [35] Joseph Redmon and Ali Farhadi. 2017. YOLO9000: Better, Faster, Stronger. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017*. IEEE Computer Society, 6517–6525. <https://doi.org/10.1109/CVPR.2017.690>

- [36] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7–12, 2015, Montreal, Quebec, Canada*, Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett (Eds.), 91–99.
- [37] Hamid Rezaatoughi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian D. Reid, and Silvio Savarese. 2019. Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019*. Computer Vision Foundation / IEEE, 658–666. <https://doi.org/10.1109/CVPR.2019.00075>
- [38] Clifford Kwok-Fung So and George Baciú. 2005. Entropy-based motion extraction for motion capture animation. *Comput. Animat. Virtual Worlds* 16, 3–4 (2005), 225–235. <https://doi.org/10.1002/cav.107>
- [39] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. 2020. FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.).
- [40] Gjorgji Strezoski and Marcel Worring. 2018. OmniArt: A Large-scale Artistic Benchmark. *ACM Trans. Multim. Comput. Commun. Appl.* 14, 4 (2018), 88:1–88:21. <https://doi.org/10.1145/3273022>
- [41] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. 2019. Deep High-Resolution Representation Learning for Human Pose Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019*. Computer Vision Foundation / IEEE, 5693–5703. <https://doi.org/10.1109/CVPR.2019.00584>
- [42] Peng Tang, Chetan Ramaiah, Yan Wang, Ran Xu, and Caiming Xiong. 2021. Proposal Learning for Semi-Supervised Object Detection. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2021, Waikoloa, HI, USA, January 3–8, 2021*. IEEE, 2290–2300. <https://doi.org/10.1109/WACV48630.2021.00234>
- [43] Antti Tarvainen and Harri Valpola. 2017. Mean Teachers are Better Role Models: Weight-averaged Consistency Targets Improve Semi-supervised Deep Learning Results. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Workshop Track Proceedings*.
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4–9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.), 5998–6008.
- [45] Karl von Amira. 1905. Die Handgebärden in den Bilderhandschriften des Sachsenspiegels. In *Abhandlungen der Bayerischen Akademie der Wissenschaften. Philosophisch-Philologische und Historische Klasse*. Franz, 162–264.
- [46] Keze Wang, Xiaopeng Yan, Dongyu Zhang, Lei Zhang, and Liang Lin. 2018. Towards Human-Machine Cooperation: Self-Supervised Sample Mining for Object Detection. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018*. Computer Vision Foundation / IEEE Computer Society, 1605–1613. <https://doi.org/10.1109/CVPR.2018.00173>
- [47] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. 2021. PVTv2: Improved Baselines with Pyramid Vision Transformer. (2021). [arXiv:2106.13797](https://arxiv.org/abs/2106.13797)
- [48] Aby Warburg. 1998. Der Eintritt des antikisierenden Idealstils in die Malerei der Frührenaissance. In *Die Erneuerung der heidnischen Antike. Kulturwissenschaftliche Beiträge zur Geschichte der europäischen Renaissance. Gesammelte Schriften*, Horst Bredekamp and Michael Diers (Eds.). Akademie Verlag, 173–176.
- [49] Aby Warburg. 1998. Dürer und die italienische Antike. In *Die Erneuerung der heidnischen Antike. Kulturwissenschaftliche Beiträge zur Geschichte der europäischen Renaissance. Gesammelte Schriften*, Horst Bredekamp and Michael Diers (Eds.). Akademie Verlag, 443–449.
- [50] Nicholas Westlake, Hongping Cai, and Peter Hall. 2016. Detecting People in Artwork with CNNs. In *Computer Vision - ECCV 2016 Workshops - Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 9913)*, Gang Hua and Hervé Jégou (Eds.). Springer, 825–841. https://doi.org/10.1007/978-3-319-46604-0_57
- [51] Bin Xiao, Haiping Wu, and Yichen Wei. 2018. Simple Baselines for Human Pose Estimation and Tracking. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8–14, 2018, Proceedings, Part VI (Lecture Notes in Computer Science, Vol. 11210)*, Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (Eds.). Springer, 472–487. https://doi.org/10.1007/978-3-030-01231-1_29
- [52] Qizhe Xie, Minh-Thang Luong, Eduard H. Hovy, and Quoc V. Le. 2020. Self-Training With Noisy Student Improves ImageNet Classification. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020*. Computer Vision Foundation / IEEE, 10684–10695. <https://doi.org/10.1109/CVPR42600.2020.01070>
- [53] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyuan Wei, Xiang Bai, and Zicheng Liu. 2021. End-to-End Semi-Supervised Object Detection with Soft Teacher. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021*. IEEE, 3040–3049. <https://doi.org/10.1109/ICCV48922.2021.00305>
- [54] Chenyang Zhang, Christine Kaeser-Chen, Grace Vesom, Jennie Choi, Maria Kessler, and Serge J. Belongie. 2019. The iMet Collection 2019 Challenge Dataset. (2019). [arXiv:1906.00901](https://arxiv.org/abs/1906.00901)
- [55] Feng Zhang, Xiatian Zhu, Hanbin Dai, Mao Ye, and Ce Zhu. 2020. Distribution-Aware Coordinate Representation for Human Pose Estimation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020*. Computer Vision Foundation / IEEE, 7091–7100. <https://doi.org/10.1109/CVPR42600.2020.00712>