

Thinking Outside the Pool: Active Training Image Creation for Relative Attributes

Aron Yu
 University of Texas at Austin
 aron.yu@utexas.edu

Kristen Grauman
 UT Austin and Facebook AI Research
 grauman@cs.utexas.edu

Abstract

Current wisdom suggests more labeled image data is always better, and obtaining labels is the bottleneck. Yet curating a pool of sufficiently diverse and informative images is itself a challenge. In particular, training image curation is problematic for fine-grained attributes, where the subtle visual differences of interest may be rare within traditional image sources. We propose an active image generation approach to address this issue. The main idea is to jointly learn the attribute ranking task while also learning to generate novel realistic image samples that will benefit that task. We introduce an end-to-end framework that dynamically “imagines” image pairs that would confuse the current model, presents them to human annotators for labeling, then improves the predictive model with the new examples. On two datasets, we show that by thinking outside the pool of real images, our approach gains generalization accuracy on challenging fine-grained attribute comparisons.

1. Introduction

Visual recognition methods are famously data hungry. Today’s deep convolutional neural networks (CNNs) achieve excellent results on various challenging tasks, and one critical ingredient is access to large manually annotated datasets [6, 8, 18, 59]. A common paradigm has emerged where a supervised learning task is defined, the relevant images/videos are scraped from the Web, and crowdworkers are enlisted to annotate them. A practical downside to this paradigm, of course, is the burden of getting all those annotations, which can be significant.

However, setting efficiency concerns aside, we contend that the standard paradigm also quietly suffers a *curation* concern. The assumption is that the more data we gather for labeling, the better. And indeed unlabeled photos are virtually unlimited. Yet—particularly when relying on Web photos—the visual variety and information content of the images amassed for labeling eventually reach a limit. The



Figure 1: Our approach learns to *actively create* training images for fine-grained attribute comparisons, focusing attention on informative pairs of images unavailable in the real image training pool. Here we show two example pairs generated to improve learning of *open* (left) and *masculine* (right).

resulting deep network inherits this limit in terms of how well it can generalize. Importantly, while *active learning* methods can try to prioritize informative instances for labeling, the curation problem remains: existing active selection methods scan the pool of manually curated unlabeled images when choosing which ones to label [12, 54, 55, 65].

We propose to address the curation challenge with *active image generation*. The main idea is to jointly learn the target visual task while also learning to generate novel realistic image samples that, once manually labeled, will benefit that task. We focus specifically on the fine-grained relative attribute task [33, 42, 49, 50, 58]: given two images, the system must infer which one more evidently exhibits an attribute (e.g., which is more *furry*, more *red*, more *smiling*). The curation problem is particularly pronounced for fine-grained relative attributes. This is because while at test time an image pair may be arbitrarily close in an attribute, it is difficult to locate extensive training samples showing such subtle differences for all attributes [61]. Our idea is for the system to dynamically “imagine” image pairs that would confuse the current learning-to-rank model, then present them to human annotators for labeling. See Figure 1.

To this end, we introduce an end-to-end deep network consisting of an attribute-conditioned image generator, a ranking network, and a *control* network. The network is seeded with real labeled image pairs of the form “image X exhibits attribute a more than image Y ”. Then the control network learns to generate latent visual parameters to present to the image generator so as to form image pairs

that would be difficult for the ranker. Thus, rather than simply perturb existing training images (à la “fooling images” [31]), the control module learns a *function* to create novel instances, potentially improving the exploration of the relevant image space. We train the ranker and controller in an adversarial manner, and we solicit manual labels for the resulting images in batches. As a result, the set of imagined training instances continues to evolve, as does the ranker.

Our main contribution is a new approach to active learning in which samples are *newly created* so as to best augment a ranker’s training data. On two challenging datasets of Shoes [59, 61] and Faces [18, 42], we demonstrate our approach’s effectiveness. Active generation focuses attention on novel training images that rapidly improve generalization—even after all available real images and their labels are exhausted. It outperforms alternative data augmentation strategies, including prior work that generates synthetic training images in a passive manner [61].

2. Related Work

Attribute Comparisons Relative attributes [33] model visual properties in a comparative manner, and a variety of applications have been explored in online shopping [24], fashion [23], biometrics [41], and graphical design [32]. Recent work explores novel learning schemes to learn relative attributes accurately, including new deep network architectures [49, 50, 58], part discovery [42], local learning [59, 60], and multi-task approaches [29]. Our contribution is an approach to actively elicit training examples for ranking, which could facilitate training for many of the above formulations.

Image Generation Photo-realistic image generation has made steady and exciting process in the recent years, often using generative adversarial networks (GANs) [13, 20, 40, 62, 63, 67] or variational auto-encoders (VAEs) [14, 25, 56]. Conditional models can synthesize an image based on a label map [20, 67] or latent attribute label [27, 52, 56]. We integrate the Attribute2Image [56] network as the image generation module in our full pipeline, though similarly equipped conditional models could also be used. Methods to generate adversarial “fooling images” also use a network to automatically perturb an image in a targeted manner [3, 30, 31]. However, rather than using adversarial generation to understand how features influence a classifier, our goal is to synthesize the very training samples that will strengthen a learned ranker. Unlike any of the above, we create images for active query synthesis.

Learning with Synthetic Images The idea of training on synthetic images but testing on real ones has gained traction in recent years. Several approaches leverage realistic human

body rendering engines to learn 3D body pose [45, 46, 53], while others extrapolate from video [7, 22, 34] or leverage 3D scene [64] or object [36, 57] models. Refiner networks that translate between the simulated and real data can help ease the inevitable domain gap [48].

Data augmentation (used widely in practice) can be seen as a hybrid approach, where each real training sample is expanded into many samples by applying low-level label-preserving transformations, or “jitter”, like scaling, rotating, mirroring, etc. [38]. Usually the jitter is manually defined, but some recent work explores how to adversarially generate “hard” jitter for body pose estimation [37], greedily select useful transformations for image classification [35], or actively evolve 3D shapes to learn shape from shading [57]. Such methods share our motivation of generating samples where they are most needed, but for very different tasks. Furthermore, our approach aligns more with active learning than data augmentation: the new synthetic samples can be distant from available real samples, and so they are manually annotated before being used for training.

In the low-shot recognition regime, several recent methods explore creative ways to hallucinate the variability around sparse real examples [9, 15, 16, 26, 43, 61], typically leveraging the observed inter-sample transformations to guide synthesis in feature-space. Among them, most relevant to our work is our “semantic jitter” approach [61], which generates relative attribute training images by altering one latent attribute variable at a time. However, whereas semantic jitter [61] uses manually defined heuristics to sample synthetic training images, we show how to dynamically derive the images most valuable to training, via an adversarial control module learned jointly with the attribute ranker.

Active Learning Active learning has been studied for decades [44]. For visual recognition problems, *pool-based* methods are the norm: the learner scans a pool of unlabeled samples and iteratively queries for the label on one or more of them based on uncertainty or expected model influence (e.g., [12, 54, 55, 65]). Active ranking models adapt the concepts of pool-based learning to select pairs for comparison [28, 39]. Hard negative mining—often used for training object detectors [11, 47]—also focuses the learner’s attention on useful samples, though in this case from a pool of already-labeled data. Rather than display one query image to an annotator, the approach in [19] selects a sample from the pool then displays a synthesized image *spectrum* around it in order to elicit feature points likely to be near the true linear decision boundary for image classification. We do not perform pool-based active selection. Unlike any of the above, our approach *creates* image samples that (once labeled) should best help the learner, and it does so in tight coordination with the ranking algorithm.

In contrast to pool-based active learning, active *query synthesis* methods request labels on novel samples from a

given distribution [1, 2, 44, 66]. When the labeler is a person (as opposed to an oracle or experimental outcome), a well known challenge with query synthesis is that the query may turn out to be unanswerable [4]. Accordingly, there is very limited prior work attempting active query synthesis for image problems. One recent study [66] uses the traditional active learning heuristic for linear SVMs [51] to generate images for toy binary image classification tasks (e.g., two MNIST digits). To our knowledge, ours is the first approach that learns to actively generate training images jointly with the image prediction task, and we demonstrate its real impact on modern well-studied datasets with complex fashion products and face images. Furthermore, rather than sample from some given input distribution, our selection approach optimizes the latent parameters of an image pair to directly affect the current deep ranking model.

3. Approach

We propose an end-to-end framework for attribute-based image comparison through active adversarial image generation. We refer to our approach as *ATTIC*, for **Ac**Tive **Tr**aining **I**mage **C**reation.

The goal is to avoid the “streetlight effect” of traditional pool-based active learning, where one looks for more training data where one already has it.¹ Rather than limit training to manually curated real images, as would standard pool-based active learning, ATTIC synthesizes image pairs that will be difficult for the ranker as it has been trained thus far. See Figure 2.

We first define the relative attributes problem (Sec. 3.1). Then we describe the key modules in our ATTIC network (Sec. 3.2). Finally we define our training procedure and the active image creation active loop (Sec. 3.3).

3.1. Relative Attributes via Ranking

Relative attributes are valuable for inferring fine-grained differences between images in terms of describable natural language properties. They have wide application in image search, zero-shot learning, and biometrics, as discussed above. At the core of the relative attribute comparison task is a learning-to-rank problem [33, 49]. A ranking function $R_{\mathcal{A}} : X \rightarrow \mathbb{R}$ for attribute \mathcal{A} (e.g., *comfortable*, *smiling*) maps an image to a real-valued score indicating the strength of the attribute in the image.

Relative attribute rankers are trained using *pairs* of training images. Let $\mathcal{P}_{\mathcal{A}} = \{(x_i, x_j)\}$ be a set of ordered pairs, where humans perceive image x_i to have “more” of attribute \mathcal{A} than image x_j . The goal is to learn a ranking function $R_{\mathcal{A}}$ that satisfies the relative orderings as much as possible: $\forall (i, j) \in \mathcal{P}_{\mathcal{A}} : R_{\mathcal{A}}(x_i) > R_{\mathcal{A}}(x_j)$. Having




Figure 2: Schematic overview of main idea. Real images (green \times 's) are used to train a deep ranking function for the attribute (e.g., the *openness* attribute for shoes). The pool of real images consists of those that are labeled (dark \times 's) and those that are unlabeled (faded \times 's). Even with all the real images labeled, the ideal ranking function may be inadequately learned. Rather than select other manually curated images for labeling (faded green \times 's), ATTIC directly generates useful synthetic training images (red \circ 's) through an adversarial learning process. The three shoes along a path represent how ATTIC iteratively evolves the control parameters to obtain the final synthetic training image pairs (not to be confused with incrementally adding “more” of a target attribute.)

trained such a function $R_{\mathcal{A}}$ we can use it to score images for their attribute strength individually, or, similarly, to judge which among a novel pair of images (x_m, x_n) exhibits the attribute more. ATTIC trains a network for each attribute \mathcal{A} ; we drop the subscript below when not needed for clarity.

3.2. End-to-End Active Training Image Creation

Let \mathcal{P}^R be an initial set of real training image pairs used to initialize the ranker R defined above. Our goal is to improve that ranker by creating synthetic training images \mathcal{P}^S , to form a hybrid training set $\mathcal{P} = \mathcal{P}^R \cup \mathcal{P}^S$.

To this end, the proposed end-to-end ATTIC framework consists of three distinct components (Fig. 3): the ranker module, the generator module, and the control module. Our model performs end-to-end adversarial learning between the ranker and the control modules. The ranker tries to produce accurate attribute comparisons, while the control module tries to produce *control parameters*—latent image parameters—that will generate difficult image pairs to confuse the current ranker. By asking human annotators to label those confusing pairs, the ranker is actively improved. We next discuss the individual modules.

3.2.1 Ranking Module

For the ranking module in ATTIC, we employ the state-of-the-art deep DeepSTN approach [49], which integrates a Siamese embedding and a spatial transformer network

¹https://en.wikipedia.org/wiki/Streetlight_effect

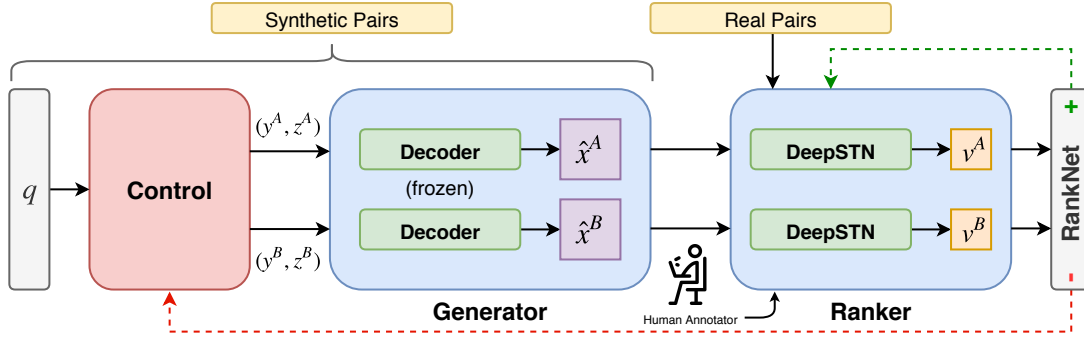


Figure 3: Architecture of our end-to-end approach consisting of three primary modules. The control module first converts the random input q into control parameters $\{(y^A, z^A), (y^B, z^B)\}$. Its architecture is detailed further in Figure 4. The generator module then generates a pair of synthetic images (\hat{x}^A, \hat{x}^B) using these control parameters. The ranker module finally uses the generated synthetic images (once manually labeled) and the real training images to train the ranking model, outputting their corresponding attribute strength (v^A, v^B) . During training, the ranking loss is fed back into the ranker (green dotted line), while the negative ranking loss is fed back into the control module (red dotted line). The decoders within the generator are pre-trained and their parameters are kept frozen throughout training.

(STN) [21] together with a RankNet ranking loss [5]. RankNet handles pairwise outputs in a single differentiable layer, using a cross-entropy loss. The rank estimates (v_i, v_j) for images (x_i, x_j) are mapped to a pairwise posterior probability using a logistic function

$$p_{ij} = \frac{1}{1 + e^{-(v_i - v_j)}}, \quad (1)$$

and the ranking loss is:

$$\mathcal{L}_{rank} = -\log(p_{ij}). \quad (2)$$

The ranker jointly learns a CNN image encoder $\phi(x)$ with a spatial transformer that localizes the informative image regions for judging a particular attribute. For example, the STN may learn to clue into the mouth region for the *smiling* attribute. The ranker combines the features of the full image and this region to compute the attribute strength:

$$v_i = R_{\mathcal{A}}(x_i) = \text{RankNet}_{\mathcal{A}}(\text{STN}(\phi(x_i))), \quad (3)$$

where $\phi(x)$ denotes application of one stack in the Siamese embedding. See [49] for details.

3.2.2 Generator Module

For the generator module, we use an existing attribute-conditioned image generator, Attribute2Image [56]. Let $\mathbf{y} \in \mathbb{R}^N$ denote an N -dimensional attribute vector containing the real-valued strength of each of the $N = |\mathcal{A}|$ attributes for an image \mathbf{x} , and let $\mathbf{z} \in \mathbb{R}^M$ denote an M -dimensional latent variable accounting for all other factors affecting the image, like lighting, pose, and background. The Attribute2Image network uses a Conditional Variational Auto-Encoder (CVAE) to learn a decoder $p_{\theta}(\hat{\mathbf{x}}|\mathbf{y}, \mathbf{z})$ that can generate realistic synthetic images $\hat{\mathbf{x}}$ conditioned

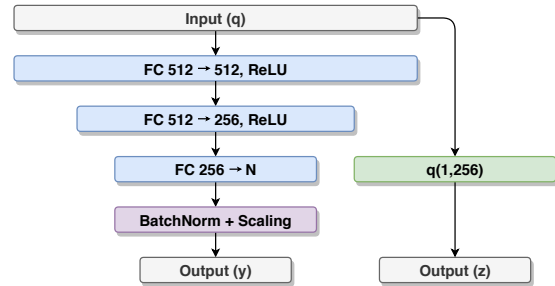


Figure 4: Architecture of the control module. The model above outputs a single set of control parameters (\mathbf{y}, \mathbf{z}) . Since we generate the synthetic images in pairs, we duplicate the architecture.

on (\mathbf{y}, \mathbf{z}) . The parameters θ are optimized by maximizing the variational lower bound of the log-likelihood $\log p_{\theta}(\hat{\mathbf{x}}|\mathbf{y})$. The network uses a multilayer perceptron to map the latent factors into their entangled hidden representation, then decodes into image pixels with a coarse-to-fine convolutional decoder. See [56] for more details.

The *attribute-conditioned* aspect of this generator allows us to iteratively refine the generated images in a semantically smooth manner, as we adversarially update its inputs (\mathbf{y}, \mathbf{z}) with the control module defined next. We pre-train the generator using $\{(x_i, \mathbf{y}_i)\}$, a disjoint set of training images labeled by their N attribute strength labels. Subsequently, we take only the decoder part of the model and use it as our generator (see Fig. 3). We freeze all parameters in the generator during active image creation, since the mapping from latent parameters to pixels is independent of the rank and control learning.

3.2.3 Control Module

As defined thus far, linking together the ranker and generator would aimlessly feed new image sample pairs to the

ranker. Next we define our control module and explain how it learns to feed pairs of intelligently chosen latent parameters to the generator for improving the ranker.

The control module is a neural network that precedes the generator (see Figure 3, left). Its input is a random seed $\mathbf{q} \in \mathbb{R}^Q$, sampled from a multivariate Gaussian. Its output is a pair of *control parameters* $\{(\mathbf{y}^A, \mathbf{z}^A), (\mathbf{y}^B, \mathbf{z}^B)\}$ for synthetic image generation. Figure 4 shows the control architecture. It is duplicated to create two branches feeding to the generator and then the Siamese network in the DeepSTN ranker.

The attribute control variable \mathbf{y} is formed by passing \mathbf{q} through a few fully-connected layers, followed by a Batch-Norm layer with scaling. In particular, we obtain the scaling parameters from the mean and standard deviation of the attribute strengths observed from the real training images, then apply them to the normalized $\mathcal{N}(0, 1)$ outputs from the BatchNorm layer. The scaling ensures that the attribute strengths are bounded within a range appropriate for the pre-trained generator.

The latent feature control variable \mathbf{z} , which captures all the non-attribute properties (e.g., pose, illumination), is sampled from a Gaussian. We simply use half of the entries from \mathbf{q} for \mathbf{z}^A and \mathbf{z}^B , respectively. This Gaussian sample agrees with the original image generator’s prior $p(\mathbf{z})$.

3.3. Training and Active Image Creation

Given the three modules, we connect them in sequence to form our active learning network. The generator and the ranker modules are duplicated for both branches to account for two images in each training pair. The decoders in the generator module are pre-trained and their parameters are kept frozen. During training, we optimize the RankNet loss for the ranker module, while at the same time optimizing the negative RankNet loss for the control module:

$$\mathcal{L}_{control} = -\mathcal{L}_{rank}. \tag{4}$$

The control module thus learns to produce parameters that generate image pairs that are difficult for the ranker to predict. This instills an adversarial effect where the control module and the ranker module are competing against each other iteratively during training. The learning terminates when the ranker converges. Our adversarial losses promote “hard” examples, and are not to be confused with adversarial “fooling” images [3, 30, 31]; image pairs generated by our method are only fed to the learner if humans can label them confidently.

Figure 5 shows examples of the *progression* of synthetic images over training iterations. ATTIC captures the interaction between the attributes and traverses the multi-attribute space to places not occupied by real training images.

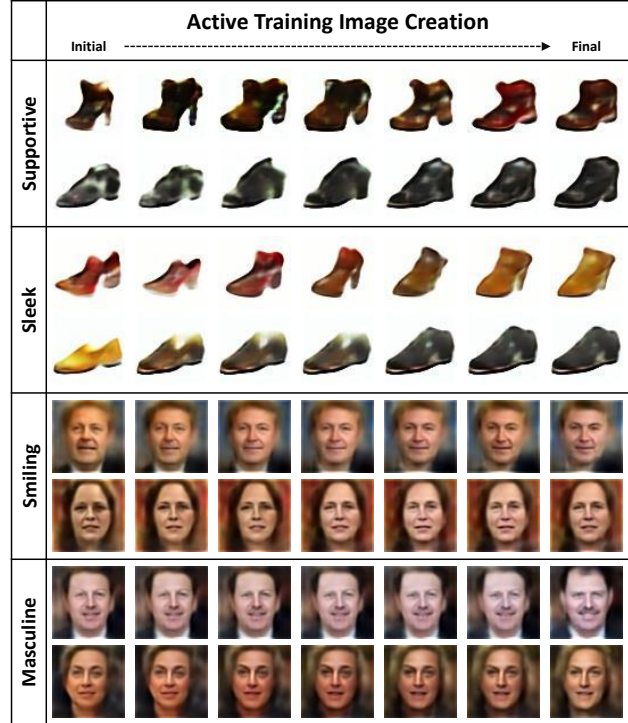


Figure 5: Visualization of the *progression* of synthetic image pairs ($\hat{\mathbf{x}}^A, \hat{\mathbf{x}}^B$) during training (not to be confused with a *spectrum* going from less to more of a given attribute [61]). Our model learns patterns between all the attributes, modifying multiple attributes simultaneously. For example, while modifying the faces for the attribute *masculine* (last row), our model learned to change the attribute *smiling* as well. The rightmost images, i.e., end of the progression, are manually labeled and augment the training data.

To generate a batch² of synthetic image pairs $\mathcal{P}^S = \{\hat{\mathbf{x}}^A, \hat{\mathbf{x}}^B\}_{i=1}^T$, we sample T vectors \mathbf{q} and push them through the control and generator. Then the batch is labeled by annotators, who judge which image shows the attribute more. The resulting pairs accepted by annotators as valid are added to the hybrid training set \mathcal{P} . We consider with both single-batch, where all synthetic pairs are generated in one go (Sec. 4.2), and multi-batch settings, where we reset the ranker module for each subsequent batch (Sec. 4.3).

The primary novelty of our approach comes from the generation of synthetic image pairs through active query synthesis. From an active learning perspective, instead of selecting more real image pairs to be labeled based on existing pool-based strategies, our approach aims to directly generate the most beneficial synthetic image pairs (Fig. 2). Furthermore, instead of choosing $(\mathbf{y}_i, \mathbf{z}_i)$ manually when generating the synthetic image pairs [61], ATTIC automates this selection in a data-driven manner.

²Here we mean “batch” in the active learning sense: a batch of examples to be manually labeled and then used to augment the training set. This is not to be confused with (*mini*)-batches for training the neural nets.

4. Experiments

To validate our active generation approach, we explore two fine-grained attribute domains: shoes and faces. To our knowledge, these are the only existing datasets with sufficient *instance-level* relative attribute labels.

4.1. Experimental Setup

Datasets We use two publicly available datasets, UT-Zappos50K [59] and LFW [18]. For each, our method uses real images to initialize training and then creates its own synthetic images.

- **Catalog Shoe Images:** We use the **UT-Zappos50K** dataset [59] with the fine-grained attributes from [61]. The dataset consists of over 50,000 shoe images from Zappos.com. There are 10 attributes (*comfort, casual, simple, sporty, colorful, durable, supportive, bold, sleek, and open*), each with about 4,000 labeled pairs.
- **Human Face Images:** We use the **LFW** dataset [18] and the **LFW-10** dataset [42]. LFW consists of over 13,000 celebrity face images from the wild along with real-valued labels on 73 attributes. LFW-10 consists of a subset of 2,000 images from LFW along with relative labels on 10 attributes. We use the 8 attributes (*bald, dark hair, big eyes, masculine, mouth open, smiling, visible forehead, and young*) in the intersection of these two datasets. For the real image pairs, there are about 500 labeled pairs per attribute from LFW-10.

All images for all methods are resized to 64×64 pixels to match the output resolution of the image generator. For all experiments, we only use high quality (high agreement/high confidence) relative labels, following the protocol of [61]. We collect annotations for our method’s automatically generated training pairs using Mechanical Turk; we obtain 5 worker responses per label and take the majority vote. Workers are free to vote for discarding a pair if they find it illegible, which happened for just 17% of the generated pairs. Of the other 83% accepted pairs, at least 4 of 5 annotators agree on the same label 63% of the time, and they rate 48% of their annotations as “very confident” and 95% as at least “somewhat confident”. See Supp File for the data collection interface and instructions to annotators.

Implementation Details During training, we validate all hyperparameters (learning rate, learning rate decay, weight decay) on a validation set. We run all experiments (including individual batches) to convergence or to a max of 250 and 100 epochs for shoes and faces, respectively.

For the individual modules, implementation details are as follows. Ranker: We pre-train the DeepSTN ranking network without the global image channel using only the real

image pairs (see [49] for details on the two rounds of training). Generator: We use the code provided by the authors of Attribute2Image [56] with all default parameters. We pre-train the image generators using a disjoint set of real images (38,000 and 11,000 images for shoes and faces, respectively), which do not have any associated relative labels. We freeze the decoder in ATTIC throughout end-to-end training for the ranker and control. Control: We initialize the layers using ReLU initialization [17].

Baselines We consider the following baselines:

- **SemJitter** [61], an existing approach for data augmentation that uses synthetic images altered by one attribute at a time using a manually defined sampling policy. We use the authors’ provided synthetic shoe image pairs, Zap50K-Synth, for this purpose. It contains 2,000 labeled pairs per attribute. For the synthetic face image pairs, we collect relative labels on 1,000 pairs per attribute.
- **Jitter**, the traditional data augmentation process where the real images are jittered through low-level geometric and photometric transformations. We follow the jitter protocol defined in [10], which includes translation, scaling, rotation, contrast, and color. The jittered image pairs retain the corresponding real pairs’ respective labels.
- **Real**, which trains with only real labeled image pairs.

Note that all methods use the same state-of-the-art ranking network for training and predictions, hence any differences in results will be attributable to the training data and augmentation strategy.

4.2. Competing with Real Training Images

First we test our hypothesis that restricting models to only available Web photos for training can be limiting. We compare the data augmentation baselines and ATTIC to the Real baseline, where all methods are given the exact same amount of total manual annotations. The Real and Jitter baselines use all n available real labeled image pairs. SemJitter and ATTIC use *half* of the real labeled image pairs ($\frac{n}{2}$), which are sampled uniformly at random, and then augment those pairs with $\frac{n}{2}$ manually labeled synthetic image pairs that they generate.

Tables 1 and 2 show the results for the Shoes and Faces datasets, respectively. We first look at the standard scenario (middle rows) where the synthetic images generated by SemJitter and our method are labeled by human annotators. Though using exactly the same amount of manual labels as the Real baseline, our method nearly always outperforms it. This shows that simply having more real image pairs labeled is not always enough; our generated sam-

	Comfort	Casual	Simple	Sporty	Colorful	Durable	Supportive	Bold	Sleek	Open	Mean
Real [49]	84.26	88.58	88.99	88.18	94.10	82.83	83.96	88.25	84.35	83.87	86.74
Jitter	82.87	87.96	86.65	87.58	93.91	80.51	84.63	88.66	83.37	79.84	85.60
SemJitter [61]	83.10	89.20	88.76	88.49	94.10	82.37	85.08	89.07	86.31	82.26	86.87
ATTIC (ours)	83.80	89.51	89.23	88.18	94.10	85.85	86.41	88.04	87.78	83.07	87.62
SemJitter-Auto [61]	84.72	88.89	89.70	89.39	94.29	83.99	86.41	89.07	85.82	83.87	87.59
ATTIC-Auto (ours)	87.04	89.20	91.57	91.21	94.48	87.94	87.31	89.90	86.31	85.75	89.07

Table 1: Accuracy for the 10 attributes in the Shoes dataset. The standard scenario (middle rows) is where the synthetic images generated by SemJitter and our method are labeled by human annotators. The “Auto” scenario (last two rows) is where an additional n unlabeled synthetic image pairs are added for all methods except Real; those images adopt their inferred attribute labels. Note that Jitter already uses all n labeled real images plus n (not manually labeled) jittered images. In all cases, all methods use exactly n total labels.

	Bald	DarkHair	BigEyes	Masculine	MouthOpen	Smiling	Forehead	Young	Mean
Real [49]	79.80	86.77	78.18	92.96	87.50	74.44	80.00	78.76	82.30
Jitter	81.82	86.87	80.00	91.55	84.17	74.44	85.26	79.79	82.99
SemJitter [61]	81.82	87.50	83.64	92.96	88.33	79.70	83.16	81.35	84.81
ATTIC (ours)	84.85	88.24	85.46	95.78	79.17	81.96	84.21	80.31	84.99
SemJitter-Auto [61]	82.83	88.24	81.82	94.37	85.00	75.19	83.16	79.28	83.73
ATTIC-Auto (ours)	85.86	90.44	83.36	94.37	82.50	76.69	85.26	78.24	84.59

Table 2: Accuracy for the 8 attributes in the Faces dataset. Format is the same as Table 1.

ples improve training across the variety of attributes in ways the existing real image pairs could not. We also tried to improve the Real baseline by adding pseudo labeled pairs bootstrapped from the attribute strengths of the images used to train the image generator, but it actually hurt the baseline slightly, likely because those pseudo labels are noisy.

Our approach also outperforms (or matches) SemJitter in 8 out of 10 shoe attributes and 6 out of 8 face attributes, with gains over 3% in some cases. This demonstrates our key advantage over SemJitter [61], which is to actively adapt the generated images to best suit the learning of the model, as opposed to what looks the best to human eyes. Unlike [61], which manually modifies one attribute at a time, our approach can modify multiple attributes simultaneously in a dynamic manner, accounting for their dependencies.

We also consider an “Auto” scenario where instead of adding $\frac{n}{2}$ generated images with their manual annotations, we start our model with all n real labeled image pairs. Then, we generate another n synthetic images and—rather than get them labeled—simply adopt their inferred attribute comparison labels. In this case, the “ground truth” ordering for attribute j for generated images \hat{x}^A and \hat{x}^B is automatically determined by the magnitudes of their associated parameter values $\mathbf{y}^A(j)$ and $\mathbf{y}^B(j)$ output by the control module. Analogously, Jitter adopts the label of the source pair it jittered. Once again, all methods use the exact same number of human-provided labels.

Tables 1 and 2 (bottom two rows) show the results. Our model performs even a bit better in this setting, suggesting that the inferred labels are often accurate, and the extra volume of “free” training pairs is helpful. We outperform (or match) SemJitter in all 10 shoe attributes and 6 out of




Figure 6: Active learning curves for the Shoes (left) and Faces (right) datasets. We show the average gain over the Real baseline after each batch of additional generated image pairs. Our approach nearly doubles the gain achieved by SemJitter for both domains.

8 face attributes. Jitter offers a slight performance boost sometimes, but can even be detrimental on these datasets. While our method performs well overall, for a couple of attributes (i.e. *mouth-open*, *young*) we underperform both Real and SemJitter. Upon inspection, we find our weaker performance there is likely due to poorer quality image generation for those attributes.

4.3. Active vs. Passive Training Image Generation

Next we examine more closely ATTIC’s active learning behavior. In this scenario, we suppose the methods have exhausted all available real training data (i.e., we use all n real labeled image pairs to initialize the model), and our goal is to augment this set. We generate the synthetic (labeled) image pairs in batches (again, not to be confused with the mini-batches when training neural networks). After each batch, we have them annotated, update the ranker’s training set, and re-evaluate it on the test set. The weights of the

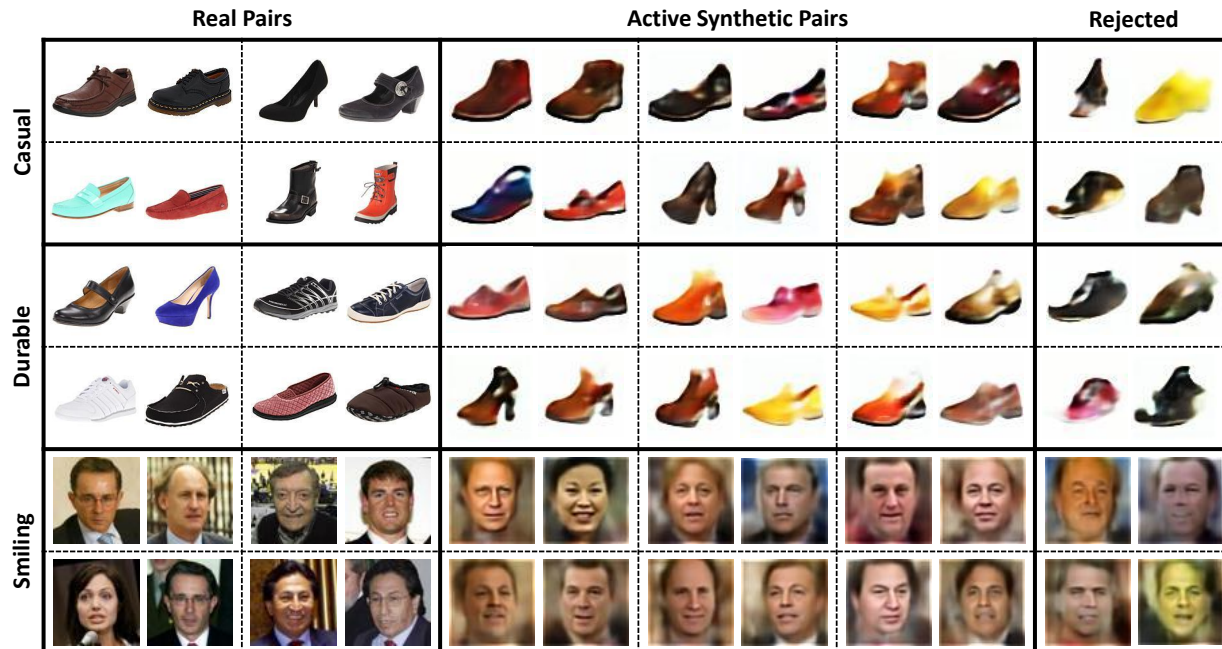


Figure 7: Sample training image pairs. Left: “Harder” real pairs that are incorrectly predicted by the baseline model. Middle: Synthetic image pairs generated by our active approach. Right: Synthetic image pairs that are rejected by the human annotators as illegible.

control module are carried over from batch to batch, while the ranker module restarts at its pre-trained state at the beginning of each batch.

Figure 6 shows the results for both datasets. We plot active learning curves to show the accuracy improvements as a function of annotator effort—steeper curves are better, as they mean the system gets more accurate with less manual labeling. We see the average gains of our approach over the Real baseline increase most sharply compared to the baselines. Our approach achieves a gain of over 3% and 8% for the two domains, respectively, which is almost double that of the SemJitter baseline. Jitter falls short once again, suggesting that traditional low-level jittering has limited impact in these fine-grained ranking tasks. Please see Supp File for the individual performance plots of each attribute. ATTIC improves more than ATTIC-Auto, though has the burden of gathering labels from annotators.

4.4. Qualitative Analysis

The results show the synthetic image pairs generated by our approach outperform those selected by the heuristic and passive selection processes of SemJitter and Jitter in most scenarios. The advantage of our active generation approach is its ability to modify the generated image pairs in a way that is best for the learning of the model.

Figure 5 shows examples of how the synthetic images look between the first and the last epoch of the training. The final pairs selected for labeling also demonstrate subtler visual differences than the initial pairs, suggesting that our

model has indeed learned to generate “harder” pairs.

Figure 7 compares these “harder” pairs to those from the real image pairs. Overall we see that the actively generated synthetic pairs tend to have fine-grained differences and/or offer visual diversity from the real training samples. The righthand side of Figure 7 shows examples of generated pairs rejected by annotators as illegible, which occurs 17% of the time. The relative low rate of rejection is an encouraging sign for making active query synthesis viable for image labeling tasks.

5. Conclusion

We introduced an approach for actively generating training image pairs that can benefit a fine-grained attribute ranker. Our approach lets the system think “outside of the pool” in annotation collection, imagining its own training samples and querying human annotators for their labels. On two difficult datasets we showed that the approach offers real payoff in accuracy for distinguishing subtle attribute differences, with consistent improvements over existing data augmentation techniques. In future work we plan to explore joint multi-attribute models for the ranker and consider how human-in-the-loop systems like ours might allow simultaneous refinement of the image generation scheme.

Acknowledgements We thank Qiang Liu, Bo Xiong, and Xinchen Yan for helpful discussions. This research is supported in part by ONR PECASE N00014-15-1-2291, NSF IIS-1514118, and a gift from Qualcomm.

References

- [1] I. Alabdulmohsin, X. Gao, and X. Zhang. Efficient active learning of halfspaces via query synthesis. In *AAAI*, 2015. 3
- [2] D. Angluin. Queries and concept learning. *Machine learning*, 1988. 3
- [3] Shumeet Baluja and Ian Fischer. Adversarial transformation networks: Learning to generate adversarial examples. Technical Report arXiv:1703.09387, 2017. 2, 5
- [4] E. Baum and K. Lang. Query learning can work poorly when a human oracle is used. In *IJCNN*, 1992. 3
- [5] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender. Learning to rank using gradient descent. In *Proceedings of International Conference on Machine Learning (ICML)*, 2005. 4
- [6] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 1
- [7] C.-Y. Chen and K. Grauman. Watching unlabeled video helps learn new human actions from very few labeled snapshots. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013 (Oral). 2
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A Large-Scale Hierarchical Image Database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 1
- [9] M. Dixit, R. Kwitt, M. Niethammer, and N. Vasconcelos. AGA: Attribute-Guided Augmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [10] A. Dosovitskiy, J. Springenberg, M. Riedmiller, and T. Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *NIPS*, 2014. 6
- [11] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 32(9), September 2010. 2
- [12] A. Freytag, E. Rodner, and J. Denzler. Selecting influential examples: active learning with expected model output changes. In *ECCV*, 2014. 1, 2
- [13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [14] K. Gregor, I. Danihelka, A. Graves, and D. Wierstra. Draw: A recurrent neural network for image generation. In *ICML*, 2015. 2
- [15] B. Hariharan and R. Girshick. Low-Shot Visual Recognition by Shrinking and Hallucinating Features. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [16] S. Hauberg, O. Freifeld, A. Larsen, J. Fisher, and L. Hansen. Dreaming more data: Class-independent distributions over diffeomorphisms for learned data augmentation. In *AISTATS*, 2016. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 6
- [18] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007. 1, 2, 6
- [19] Miriam W Huijser and Jan C van Gemert. Active decision boundary annotation with deep generative models. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [20] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 2
- [21] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. Spatial transformer networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2015. 4
- [22] A. Khoreva, R. Benenson, E. Ilg, T. Brox, and B. Schiele. Lucid data dreaming for object tracking. Technical Report arXiv:1703.09554, 2017. 2
- [23] M. Kiapour, K. Yamaguchi, A. Berg, and T. Berg. Hipster wars: Discovering elements of fashion styles. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2014. 2
- [24] A. Kovashka, D. Parikh, and K. Grauman. WhittleSearch: Interactive image search with relative attribute feedback. *International Journal of Computer Vision (IJCV)*, 115(2):185–210, Nov 2015. 2
- [25] T. Kulkarni, W. Whitney, P. Kohli, and J. Tenenbaum. Deep convolutional inverse graphics network. In *NIPS*, 2015. 2
- [26] R. Kwitt, S. Hegenbart, and M. Niethammer. One-Shot Learning of Scene Locations via Feature Trajectory Transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [27] Guillaume Lample, Neil Zeghidour, Nicolas Usunier, Antoine Bordes, Ludovic DENOYER, et al. Fader networks: Manipulating images by sliding attributes. In *NIPS*, pages 5963–5972, 2017. 2
- [28] L. Liang and K. Grauman. Beyond comparing image pairs: Setwise active learning for relative attributes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2
- [29] Zihang Meng, Nagesh Adluru, Hyunwoo J. Kim, Glenn Fung, and Vikas Singh. Efficient relative attribute learning using graph neural networks. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2018. 2
- [30] S. Moosavi-Dezfooli, A. Fawzi, and P. Frossard. DeepFool: A Simple and Accurate Method to Fool Deep Neural Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 5
- [31] A. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 5
- [32] P. O’Donovan, J. Libeks, A. Agarwala, and A. Hertzmann. Exploratory font selection using crowdsourced attributes. In *SIGGRAPH*, 2014. 2

- [33] D. Parikh and K. Grauman. Relative Attributes. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2011. 1, 2, 3
- [34] D. Park and D. Ramanan. Articulated pose estimation with tiny synthetic videos. In *ChaLearn Workshop, CVPR*, 2015. 2
- [35] Mattis Paulin, Jrme Revaud, Zaid Harchaoui, Florent Perronnin, and Cordelia Schmid. Transformation pursuit for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2
- [36] X. Peng, B. Sun, K. Ali, and K. Saenko. Learning deep object detectors from 3d models. In *ICCV*, 2015. 2
- [37] Xi Peng, Zhiqiang Tang, Fei Yang, Rogerio S Feris, and Dimitris Metaxas. Jointly optimize data augmentation and network training: Adversarial data augmentation in human pose estimation. In *CVPR*, 2018. 2
- [38] L. Perez and J. Wang. The effectiveness of data augmentation in image classification using deep learning. Technical Report arXiv:1712.04621, 2017. 2
- [39] Buyue Qian, Xiang Wang, Fei Wang, Hongfei Li, Jieping Ye, and Ian Davidson. Active learning from relative queries. In *IJCAI International Joint Conference on Artificial Intelligence*, 2013. 2
- [40] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016. 2
- [41] D. Reid and M. Nixon. Human identification using facial comparative descriptions. In *ICB*, 2013. 2
- [42] R. Sandeep, Y. Verma, and C. Jawahar. Relative parts: Distinctive parts for learning relative attributes. In *CVPR*, 2014. 1, 2, 6
- [43] Eli Schwartz, Leonid Karlinsky, Joseph Shtok, Sivan Harary, Mattias Marder, Rogerio Feris, Abhishek Kumar, Raja Giryes, and Alex M. Bronstein. Delta-encoder: An effective sample synthesis method for few-shot object recognition. Technical Report arXiv:1806.04734, 2018. 2
- [44] Burr Settles. Active learning literature survey. Technical report, 2010. 2, 3
- [45] G. Shakhnarovich, P. Viola, and T. Darrell. Fast Pose Estimation with Parameter-Sensitive Hashing. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2003. 2
- [46] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake. Real-time human pose recognition in parts from single depth images. In *CVPR*, 2011. 2
- [47] A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016. 2
- [48] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from Simulated and Unsupervised Images through Adversarial Training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [49] K. Singh and Y. J. Lee. End-to-end localization and ranking for relative attributes. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2016. 1, 2, 3, 4, 6, 7
- [50] Y. Souri, E. Noury, and E. Adeli. Deep relative attributes. In *ACCV*, 2016. 1, 2
- [51] S. Tong and D. Koller. Support vector machine active learning with applications to text classification. In *ICML*, 1998. 3
- [52] P. Upchurch, J. Gardner G. Pleiss, R. Pless, N. Snavely, K. Bala, and K. Weinberger. Deep feature interpolation for image content changes. In *CVPR*, 2017. 2
- [53] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *CVPR*, 2017. 2
- [54] S. Vijayanarasimhan and K. Grauman. What’s it going to cost you?: Predicting effort vs. informativeness for multi-label image annotations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 1, 2
- [55] S. Vijayanarasimhan and K. Grauman. Large-scale live active learning: Training object detectors with crawled data and crowds. *International Journal of Computer Vision (IJCV)*, 108(1):97–114, May 2014. 1, 2
- [56] X. Yan, J. Yang, K. Sohn, and H. Lee. Attribute2Image: Conditional image generation from visual attributes. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2016. 2, 4, 6
- [57] D. Yang and J. Deng. Shape from shading through shape evolution. Technical Report arXiv:1712.02961, 2017. 2
- [58] X. Yang, T. Zhang, C. Xu, S. Yan, M. Hossain, and A. Ghoneim. Deep relative attributes. *IEEE Trans. on Multimedia*, 18(9), Sept 2016. 1, 2
- [59] A. Yu and K. Grauman. Fine-Grained Visual Comparisons with Local Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 1, 2, 6
- [60] A. Yu and K. Grauman. Just Noticeable Differences in Visual Attributes. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 2
- [61] A. Yu and K. Grauman. Semantic Jitter: Dense Supervision for Visual Comparisons via Synthetic Images. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2, 5, 6, 7
- [62] Gang Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Generative adversarial network with spatial attention for face attribute editing. In *ECCV*, 2018. 2
- [63] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, 2017. 2
- [64] Y. Zhang, S. Song, E. Yumer, M. Savva, J. Lee, H. Jin, and T. Funkhouser. Physically-based rendering for indoor scene understanding using convolutional neural networks. In *CVPR*, 2017. 2
- [65] L. Zhao, G. Sukthankar, and R. Sukthankar. Robust active learning using crowdsourced annotations for activity recognition. In *HCOMP*, 2011. 1, 2
- [66] Jia-Jie Zhu and Jos Bento. Generative adversarial active learning. Technical Report arXiv:1702.07956, 2017. 3

- [67] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
2