

Cross-GAN Auditing: Unsupervised Identification of Attribute Level Similarities and Differences between Pretrained Generative Models

Matthew L. Olson¹, Shusen Liu², Rushil Anirudh², Jayaraman J. Thiagarajan², Peer-Timo Bremer², and Weng-Keen Wong¹

¹ Oregon State University - EECS, ²Lawrence Livermore National Laboratory - CASC
 {olsomatt,wongwe}@oregonstate.edu, {liu42,anirudh1,jayaramanthi1,bremer5}@llnl.gov

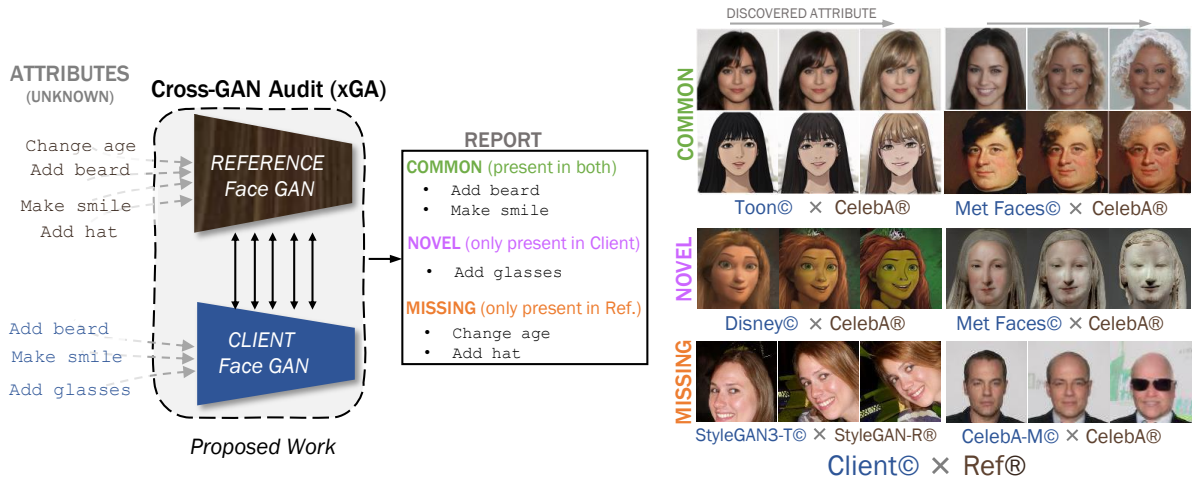


Figure 1. We introduce (xGA) an approach for fully unsupervised cross-GAN auditing and validation. Given two pre-trained GANs (Reference & Client), xGA evaluates the client by identifying three types of semantic attributes – (a) Common: those that exist in both models, (b) Novel: those only present in the client and (c) Missing: those that only exist in the reference. On the right, we show results across multiple studies, that among others include notable shifts in distribution between the Reference (CelebA) to Client (Toon, Disney, Met Faces). xGA also lends itself easily to comparing models with different properties on the same dataset as shown on the bottom right for StyleGAN3-T vs. StyleGAN-R. And CelebA-M is a control dataset we create that does not contain glasses, ties and smiles.

Abstract

Generative Adversarial Networks (GANs) are notoriously difficult to train especially for complex distributions and with limited data. This has driven the need for tools to audit trained networks in human intelligible format, for example, to identify biases or ensure fairness. Existing GAN audit tools are restricted to coarse-grained, model-data comparisons based on summary statistics such as FID or recall. In this paper, we propose an alternative approach that compares a newly developed GAN against a prior baseline. To this end, we introduce Cross-GAN Auditing (xGA) that, given an established “reference” GAN and a newly proposed “client” GAN, jointly identifies intelligible attributes that are either common across both GANs, novel to the client GAN, or missing from the client GAN.

This provides both users and model developers an intuitive assessment of similarity and differences between GANs. We introduce novel metrics to evaluate attribute-based GAN auditing approaches and use these metrics to demonstrate quantitatively that xGA outperforms baseline approaches. We also include qualitative results that illustrate the common, novel and missing attributes identified by xGA from GANs trained on a variety of image datasets¹.

1. Introduction

Generative Adversarial Networks (GANs) [12, 19–21] have become ubiquitous in a range of high impact commercial and scientific applications [6, 8–10, 13]. With this pro-

¹Source code is available at https://github.com/mattolson93/cross_gan_auditing

lific use comes a growing need for investigative tools that are able to evaluate, characterize and differentiate one GAN model from another, especially since such differences can arise from a wide range of factors – biases in training data, model architectures and hyper parameters used in training etc. In practice, this has been mostly restricted to comparing two or more GAN models against the dataset they were trained on using summary metrics such as Fréchet Inception Distance (FID) [16] and precision/recall [20] scores.

However, in many real world scenarios, different models may not even be trained on the same dataset, thereby making such summary metrics incomparable. More formally, if we define the model comparison problem as one being between a known – and presumably well vetted – *reference* GAN and a newly developed *client* GAN. For example, the reference GANs can correspond to models purchased from public market places such as AWS [1], Azure [3], or GCP [2], or to community-wide standards. Furthermore, there is a critical need for more fine-grained, interpretable, investigative tools in the context of fairness and accountability. Broadly, these class of methods can be studied under the umbrella of AI model *auditing* [4, 7, 32]. Here, the interpretability is used in the context to indicate that the proposed auditing result will involve human intelligible attributes, rather than summary statistics that do not have explicit association with meaningful semantics.

While auditing classifiers has received much attention in the past [32], GAN auditing is still a relatively new research problem with existing efforts focusing on model-data comparisons, such as identifying how faithfully a GAN recovers the original data distribution [4]. In contrast, we are interested in developing a more general framework that enables a user to visually audit a “client” GAN model with respect to the “reference”. This framework is expected to support different kinds of auditing tasks: (i) comparing different GAN models trained on the same dataset (e.g. StyleGAN3-Rotation and StyleGAN3-Translate on FFHQ); (ii) comparing models trained on datasets with different biases (e.g., StyleGAN with race imbalance vs StyleGAN with age imbalance); and finally (iii) comparing models trained using datasets that contain challenging distribution shifts (e.g., CelebA vs Toons). Since these tools are primarily intended for human experts and auditors, interpretability is critical. Hence, it is natural to perform auditing in terms of human intelligible attributes. Though there has been encouraging progress in automatically discovering such attributes from a single GAN in the recent years [14, 28, 39, 40, 43] they are not applicable to our setting with multiple GANs.

Proposed work We introduce cross-GAN auditing (xGA), an unsupervised approach for identifying attribute similarities and differences between client GANs and reference models (which could be pre-trained and potentially unrelated). Since the GANs are trained independently, their la-

tent spaces are disparate and encode different attributes, and thus they are not directly comparable. Consequently, discovering attributes is only one part of the solution; we also need to ‘align’ humanly meaningful and commonly occurring attributes across the individual latent spaces.

Our audit identifies three distinct sets of attributes: (a) common: attributes that exist in both client and reference models; (b) novel: attributes encoded only in the client model; (c) missing: attributes present only in the reference. In order to identify common attributes, xGA exploits the fact that shared attributes should induce similar changes in the resulting images across both the models. On the other hand, to discover novel/missing attributes, xGA leverages the key insight that attribute manipulations unique to one GAN can be viewed as out of distribution (OOD) to the other GAN. Using empirical studies with a variety of StyleGAN models and benchmark datasets, we demonstrate that xGA is effective in providing a fine-grained characterization of generative models.

Contributions (i) We present the first cross-GAN auditing framework that uses a unified, attribute-centric method to automatically discover common, novel, and missing attributes from two or more GANs; (ii) Using an external, robust feature space for optimization, xGA produces high-quality attributes and achieves effective alignment even across challenging distribution shifts; (iii) We introduce novel metrics to evaluate attribute-based GAN auditing approaches; and (iv) We evaluate xGA using StyleGANs trained on CelebA, AFHQ, FFHQ, Toons, Disney and MetFaces, and also provide a suite of controlled experiments to evaluate cross-GAN auditing methods.

2. Related Work

Attribute Discovery Several approaches have been successful in extracting attribute directions in StyleGAN’s latent space in the past few years. InterfaceGAN [33] used an external classifier and human annotations to label sampled images in order to build a simple linear model that captures the attribute direction in a GAN’s latent space. GANSpace [14] applies PCA to these intermediate representations to find the large factors of variation and then re-projects these directions onto a GAN’s latent space. Similarly, SeFa [34] directly captures these directions via matrix factorization of the affine mapping weights in styleGAN, which identify directions of large changes without the need to sample the latent space. An alternative strategy is to directly learn the interpretable directions through a jointly-trained predictive model by assuming that the more predictive variations are more likely to be semantically meaningful [39]– or that using a Hessian penalty [28], or Jacobian [40], in the image space enables learning of directions. LatentCLR [43] used a similar optimization framework, but instead of training a separate predictive model, it leveraged

the GAN’s internal representation and adopted a contrastive loss [11] for attribute discovery.

Model Auditing With increased awareness of the societal impact of machine learning models, there is an increased interest in characterizing and criticizing model behavior under the broad umbrella of auditing [32, 42]. There has been relatively less work in auditing generative models. For example, [4] introduce a new performance metric for generative models that measures fidelity, diversity, and generalization. Another related work is from Bau et al., [7] who investigate what a GAN cannot generate, whereas our interest is in distinguishing a client GAN from a reference GAN.

Interpretation of Domain Shift Some of the most related work comes from methods that aim for characterizing domain shift [26, 27], but these methods are limited to specific settings: either relying on human intervention [26] or needing a disentangled generator in the input [27]. An indirect way to obtain aligned attributes is via *aligned GANs*—GANs where one is fine-tuned from the other [41], [29]. Here, the attribute direction will be inherent to the children models, eliminating the need to do joint discovery to identify similar attributes. However, obtaining an *aligned GAN* through a separate fine-tuning process for attribute discovery across distributions is neither practical nor feasible.

3. Methods

We approach GAN auditing as performing attribute-level comparison to a reference GAN. For simplicity, we consider the setup where there is a single reference and client model to perform auditing, though xGA can be used even with multiple reference or client models (see experiments). Let us define the reference and client generators as $\mathcal{G}_r : \mathcal{Z}_r \mapsto \mathcal{X}_r$ and $\mathcal{G}_c : \mathcal{Z}_c \mapsto \mathcal{X}_c$ respectively. Here, \mathcal{Z}_r and \mathcal{Z}_c refer to the corresponding latent spaces and the generators are trained to approximate the data distributions $P_r(x)$ and $P_c(x)$. Our formulation encompasses the scenario where $P_r(x) = P_c(x)$ but the model architectures are different, or the challenging setting of $P_r(x) \neq P_c(x)$ (e.g., CelebA faces vs Met Faces datasets).

The key idea of xGA is to audit a client model \mathcal{G}_c via attribute (i.e., directions in the latent space) comparison to a reference model, in lieu of computing summary scores (e.g., FID, recall) from the synthesized images. In order to enable a fine-grained, yet interpretable, analysis of GANs, xGA performs automatic discovery and categorization of latent attributes: (i) *common*: attributes that are shared by both the models; (ii) *missing*: attributes that are captured by \mathcal{G}_r , but not \mathcal{G}_c ; (iii) *novel*: attributes that are encoded in \mathcal{G}_c but not observed in the reference model. We express this new categorization scheme in figure 2. Together, these latent attributes can provide a holistic characterization of GANs, while circumventing the need for customized metrics or human-centric analysis.

	\bar{h}_i^1	\bar{h}_j^1	h_i^1	h_j^1	\bar{h}_i^n	\bar{h}_j^n	h_i^n	h_j^n
\bar{h}_i^1	+	+	+	-	-	-	-	-
\bar{h}_j^1	+	+	+	-	-	-	-	-
h_i^1	+	+	+	-	-	-	-	-
h_j^1	+	+	+	-	-	-	-	-
\bar{h}_i^n	-	-	-	-	+	-	-	-
\bar{h}_j^n	-	-	-	-	+	-	-	-
h_i^n	-	-	-	-	-	-	+	-
h_j^n	-	-	-	-	-	-	+	+

Traditional Contrastive Loss: '+' Positive, '-' Negative

xGA: '+' Common, '-' Missing, 'x' Novel

Figure 2. A table showing the proposed xGA modifications to typical contrastive loss with a simple two attribute model.

Latent attributes: Following state-of-the-art approaches such as LatentCLR [43], we define attributes as direction vectors in the latent space of a GAN. For any sample $z \in \mathcal{Z}_c$ and a direction vector δ_n , we can induce attribute-specific manipulation to the corresponding image as

$$\mathcal{D} : (z, \delta_n) \rightarrow z + \alpha \delta_n, \text{ where } \delta_n = \frac{\mathbf{M}_n z}{\|\mathbf{M}_n z\|}, \quad (1)$$

for a scalar α , and a learnable matrix \mathbf{M}_n . In other words, we consider the attribute change to be a linear model defined by the learnable direction δ_n . The manipulated image can then be obtained as $\mathcal{G}_c(\mathcal{D}(z, \delta_n))$, or in shorter notation $\mathcal{G}_c(z, \delta_n)$. Note that these latent attributes are not pre-specified and are discovered as part of the auditing process.

3.1. Common Attribute Discovery

Identifying common attributes between the client and reference GAN models is challenging, since it requires that the latent directions are *aligned*, i.e., the exact same semantic change must be induced in unrelated latent spaces. When distilling from a parent model, i.e., training Toons from Faces, attributes appear to align naturally, even under severe distribution shifts [41]. However, this does not hold true when the two models are trained independently, which requires us to solve the joint problem of identifying the attributes as well as explicitly aligning them.

Formally, for a common attribute, we want the semantic change (in the generated images) induced by manipulating any sample $z \in \mathcal{Z}_c$ along a direction δ in the client GAN’s latent space to match the change in the direction $\bar{\delta}$ from the reference GAN’s latent space for any $\bar{z} \in \mathcal{Z}_r$. In other words, $S(\mathcal{G}_c(z, \delta), \mathcal{G}_c(z)) \approx S(\mathcal{G}_r(\bar{z}, \bar{\delta}), \mathcal{G}_r(\bar{z})), \forall z \in \mathcal{Z}_c, \bar{z} \in \mathcal{Z}_r$. Here, S denotes an *oracle* detector (e.g., human subject test) which measures the semantic changes between the original sample and that obtained by manipulating the common attribute.

However, in practice, such a semantic change detector is not accessible and we need to construct a surrogate mechanism to quantify the alignment, i.e.,

$$\min_{\delta_n, \bar{\delta}_n} \mathcal{L} \left(\mathcal{G}_c(z, \delta_n), \mathcal{G}_r(\bar{z}, \bar{\delta}_n) \right), \forall z \in \mathcal{Z}_c, \forall \bar{z} \in \mathcal{Z}_r, \quad (2)$$

for a common attribute pair $(\delta_n, \bar{\delta}_n)$. Any choice of the loss function \mathcal{L} must satisfy two key requirements: (a) identify high-quality, latent directions within each of the latent spaces; (b) encourage cross-GAN alignment such that similar attributes end up being strongly correlated under the loss function. For example, in the case of a single GAN, the LatentCLR [43] approach learns distinct directions using a contrastive objective that defines positive samples as those that have all been perturbed in the same direction, while manipulations in all other directions are considered negative². However, this approach is not suitable for our setting because of a key limitation – alignment requires us to operate in a common feature space so that semantics across the two models are comparable. To address this, we first modify the objective to operate in the latent space of an external, pre-trained feature extractor \mathcal{F} . In order to support alignment even in the scenario where $P_c(x) \neq P_r(x)$, we can choose \mathcal{F} that is robust to commonly occurring distributional shifts.

Our approach works on mini-batches of size B samples each, randomly drawn from \mathcal{Z}_c and \mathcal{Z}_r respectively. For the i^{th} sample in a mini-batch from \mathcal{Z}_c , let us define the vector h_i^n as the divergence between the output of the GAN before and after perturbing along the n^{th} latent direction, computed in the feature space of \mathcal{F} , i.e., $h_i^n = \mathcal{F}(\mathcal{G}_c(z_i, \delta_n)) - \mathcal{F}(\mathcal{G}_c(z_i))$. Similarly, we define the divergence $\bar{h}_j^n = \mathcal{F}(\mathcal{G}_r(\bar{z}_j, \bar{\delta}_n)) - \mathcal{F}(\mathcal{G}_r(\bar{z}_j))$ for the reference GAN. Next, we measure the semantic similarity between the divergence vectors as $g(h_i^n, \bar{h}_j^n) = \exp(\cos(h_i^n, \bar{h}_j^n)/\tau)$, where τ is the temperature parameter, and \cos refers to cosine similarity. Now, the loss function for inferring a common attribute can be written as

$$\begin{aligned} \mathcal{L}_{\text{xent}}(\delta_n, \bar{\delta}_n, \lambda_a) = & \sum_{i=1}^B \sum_{j \neq i}^B g(h_i^n, h_j^n) + g(\bar{h}_i^n, \bar{h}_j^n) + \lambda_a g(\bar{h}_i^n, h_j^n) \\ - \log & \frac{\sum_{i=1}^B \sum_{l=1}^N \mathbb{1}_{[l \neq n]} \left(g(h_i^l, h_j^n) + g(\bar{h}_i^l, \bar{h}_j^n) + g(\bar{h}_i^l, h_j^n) \right)}{\sum_{i=1}^B \sum_{l=1}^N \mathbb{1}_{[l \neq n]} \left(g(h_i^l, h_j^n) + g(\bar{h}_i^l, \bar{h}_j^n) + g(\bar{h}_i^l, h_j^n) \right)} \end{aligned} \quad (3)$$

Here N denotes the total number of attributes. While the first two terms in the numerator are aimed at identifying distinct attributes from \mathcal{G}_c and \mathcal{G}_r , the third term enforces the

²Other single GAN methods could be adapted, but LatentCLR’s flexible loss requires less computation without the need to enforce orthogonality at every learning step.

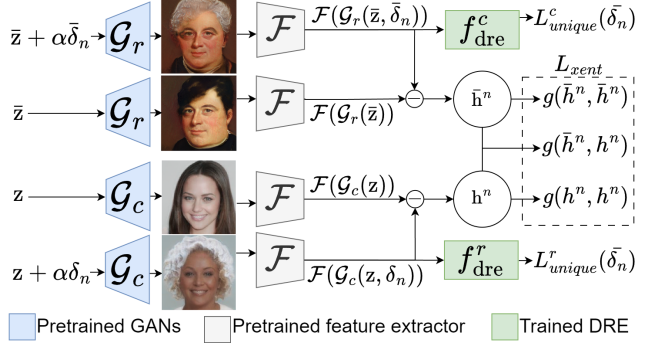


Figure 3. A diagram of our xGA model. \mathcal{G}_r , \mathcal{G}_c , and \mathcal{F} are fixed pretrained models. δ_n and $\bar{\delta}_n$ are direction models trained to learn aligned attributes between the two Generators using the features of \mathcal{F} , and f_{dre} are regularization models for unique attributes.

pair $(\delta_n, \bar{\delta}_n)$ to induce similar semantic change. When the λ_a parameter is set to 0, this optimization reinforces self-similarity of the attributes, without cross-similar semantics. The terms in the denominator are based on the negative pairs (divergences from different latent directions) to enable contrastive training.

3.2. Novel & Missing Attribute Discovery

A key component of our GAN auditing framework is the discovery of interpretable attributes that are unique to or missing from the client GAN’s latent space. This allows practitioners to understand the novelty and limitations of a GAN model with respect to a well-established reference GAN. To this end, we exploit the key intuition that images synthesized by manipulating an attribute specific to the client model can manifest as out-of-distribution (OOD) to the reference model (and vice versa).

In order to characterize the OOD nature of such realizations, we define a likelihood score in the feature space from \mathcal{F} , which indicates whether a given sample is out of distribution. More specifically, we use the Density Ratio Estimation (DRE) [25, 36] method that seeks to approximate the ratio: $\gamma(x) = \frac{P(x)}{Q(x)}$ for any sample x . When the ratio is low, it is likely that x is from the distribution Q and hence OOD to P . We choose DRE, specifically the Kullback-Liebler Importance Estimation Procedure (KLIEP) [37], over other scoring functions because it is known to be highly effective at accurately detecting outliers [24].

We pre-train two separate DRE models to approximate $\gamma_c(z)$, and $\gamma_r(\bar{z})$, wherein we treat data from $\mathcal{F}(\mathcal{G}_c(z))$ as P and $\mathcal{F}(\mathcal{G}_r(\bar{z}))$ as Q for the former, and vice versa for the latter. These DRE models are implemented as 2-layer MLP networks, $f_{\text{dre}}^c(\cdot)$, $f_{\text{dre}}^r(\cdot)$, such that

$$\hat{\gamma}_c(z) = f_{\text{dre}}^c(\mathcal{F}(\mathcal{G}_c(z))) \text{ and } \hat{\gamma}_r(\bar{z}) = f_{\text{dre}}^r(\mathcal{F}(\mathcal{G}_r(\bar{z}))), \quad (4)$$

where \mathcal{F} is the same feature extractor from (3). We pass the

output of the MLPs through a softplus ($\varphi(x) = \log(1 + e^x)$) function to ensure non-negativity. As stated previously, we use the KLIEP method to train DRE models. Using Section 4.1 of [24], the KLIEP loss used for training is defined as:

$$\mathcal{L}_{\text{KLIEP}}^c = \frac{1}{T_2} \sum_{j=1}^{T_2} \hat{\gamma}_c(\bar{z}_j) - \frac{1}{T_1} \sum_{i=1}^{T_1} \ln \hat{\gamma}_c(z_i), \quad (5)$$

where \bar{z}_j and z_i are random samples drawn from the latent spaces \mathcal{Z}_r and \mathcal{Z}_c respectively (with T_1 and T_2 total samples). Similarly, we can define the KLIEP loss term for the reference model as:

$$\mathcal{L}_{\text{KLIEP}}^r = \frac{1}{T_1} \sum_{i=1}^{T_1} \hat{\gamma}_r(z_i) - \frac{1}{T_2} \sum_{j=1}^{T_2} \ln \hat{\gamma}_r(\bar{z}_j). \quad (6)$$

We also investigated using log-loss functions to train the DRE model, but found it to be consistently inferior to the KLIEP losses (see supplement for details). Finally, we use the pre-trained DRE models from the client and reference GAN data to identify novel and missing attributes, where for a given attribute n in the reference GAN, we can enforce its uniqueness by utilizing the client DRE model to give us $\mathcal{L}_{\text{Unique}}^r(\delta_n) = \hat{\gamma}_c(z, \delta_n)$ and similarly for the client GAN we can use the reference DRE model $\mathcal{L}_{\text{Unique}}^c(\bar{\delta}_n) = \hat{\gamma}_r(\bar{z}, \bar{\delta}_n)$. Note, we interpret the novel attributes from the reference GAN as the missing attributes for the client GAN.

3.3. Overall Objective

We now present the overall objective of xGA to identify N_c common, N_n novel and N_m missing attributes simultaneously. Denoting the total number of attributes $N = N_c + \max(N_n, N_m)$, the total loss can be written as:

$$\begin{aligned} \mathcal{L}_{\text{xGA}} = & \sum_{n=1}^N \mathcal{L}_{\text{xent}}(\delta_n, \bar{\delta}_n, \mathbb{1}_{[n \leq N_c]} \lambda_a) \\ & + \lambda_b \left[\sum_{p=N_c+1}^{N_c+N_n} \mathcal{L}_{\text{Unique}}^c(\bar{\delta}_p) + \sum_{q=N_c+1}^{N_c+N_m} \mathcal{L}_{\text{Unique}}^r(\delta_q) \right] \end{aligned}$$

Here, the hyper-parameter λ_b is the penalty for enforcing the attributes between the two latent spaces to be disparate (missing/novel). And we set $g(\cdot, \cdot) = 0$ in $\mathcal{L}_{\text{xent}}$ if one of the directions vectors does not exist (i.e. when $N_n \neq N_m$).

4. Experiments

In order to systematically evaluate the efficacy of our proposed GAN audit approach, we consider a suite of GAN models trained using several benchmark datasets. In this section, we present both qualitative and quantitative assessments of xGA, and additional results are included in the Supplementary Material.

4.1. Datasets and GAN Models

For most experiments, we used a StyleGANv2 [20] trained on the CelebA [23] dataset as our reference GAN model. This choice is motivated both by its wide-spread use as well as the availability of fine-grained, ground truth attributes for each of the face images in CelebA, and to ensure that this model is fully independent from other client GANs (e.g., ToonGAN is finetuned from FFHQ GAN). In one experiment for the AFHQ dataset, we used a StyleGANv2 trained using only *cat* images from AFHQ as the reference. Also, we considered FFHQ-trained StyleGANv3 [18] and non-StyleGAN architectures such as GANformer [17] for defining the reference (see supplement).

In our empirical study, we constructed a variety of (StyleGANv2) client models and performed xGA: (i) 5 trained with different CelebA subsets constructed by excluding images specific to a chosen attribute (hat, glasses, male, female and beard); (ii) 2 trained with CelebA subsets constructed by excluding images containing any of a chosen set of attributes (beards|hats, smiles|glasses|ties); (iv) 3 transferred GANs for Met Faces, cartoons [5], and Disney images [5] respectively.

4.2. Training Settings

In all our experiments, xGA training is carried out for 10,000 iterations with random samples drawn from \mathcal{Z}_c and \mathcal{Z}_r . We fixed the desired number of attributes to be $N_c = 12$, $N_n = 4$ and $N_m = 4$. Note, this choice was to enable training xGA on a single 15GB Tesla T4 GPU. With the StyleGAN2 models, our optimization takes 4 hours; StyleGAN3 takes 12 hours due to gradient check-pointing. For all latent directions $\{\delta_n\}$ and $\{\bar{\delta}_n\}$, we set $\alpha = 3$ and this controls how far we manipulate each sample in a given direction. In each iteration, the effective batch size was 10, wherein 2 samples were used to construct a positive pair and a subset of 5 directions were randomly chosen for updating (enforced due to memory constraints). We used the Adam [22] optimizer with learning rate 0.001 to update the latent direction parameters. Note, all other model parameters (generators, feature extractor, DRE models) were fixed and never updated. Following common practice with StyleGANs, the attributes are modeled in the style space and the generator's outputs are appropriately resized to fit the size requirements of the chosen feature extractor.

For our optimization objective, we set the hyper-parameter $\lambda_a = 0.1$ in \mathcal{L}_{xGA} . To perform DRE training, we used 2-layer MLPs trained via the Adam optimizer for 1000 iterations to minimize the KLIEP losses specified in (5) and (6). At each step, we constructed batches of 32 samples from both reference and client GANs, and projected them into the feature space of \mathcal{F} . Lastly, we set $\lambda_b = 1.0$; we explore tuning this parameter in the supplement, finding it to be relatively insensitive.

4.3. Evaluation: Common Attribute Discovery

We begin by evaluating the ability of xGA in recovering common attributes across reference and client models. As mentioned earlier, for effective alignment, the choice of the feature extractor is critical. More specifically, \mathcal{F} must be sufficiently expressive to uncover aligned attributes from both client and reference models. Furthermore, it is important to handle potential distribution shifts across the datasets used to train the GAN models. Hence, a feature extractor that can be robust to commonly occurring distribution shifts is expected to achieve effective alignment via (3). In fact, we make an interesting observation that performing attribute discovery in such an external feature space leads to improved disentanglement in the inferred latent directions. For all results reported here, we used a robust variant of ResNet that was trained to be adversarially robust to style variations [35]. Please refer to the ablation in Section 4.5 for a comparison of different choices.

Qualitative results In Figure 4, we show several examples of common attributes identified by xGA for different client-reference pairs, we observe that xGA finds non-trivial attributes. For example, the “sketchify” attribute which naturally occurs in Met Faces (a dataset of paintings), is surprisingly encoded even in the reference CelebA GAN (which only consists of photos of people). We also show examples of other interesting attributes such as “orange fur” in the case of dog-GAN \times cat-GAN or “blonde hair” in the case of Disney-GAN \times CelebA-GAN. These results indicate that our proposed alignment objective, when coupled with a robust feature space, can effectively reveal common semantic directions across the client and reference models. We include several additional examples in the supplement.

Quantitative results To perform more rigorous quantitative comparisons, we setup a controlled experiment using 7 client models corresponding to different CelebA subsets (obtained by excluding images pertinent to specific characteristics). As discussed earlier, we use a standard CelebA StyleGANv2 as the common reference model across all 7 experiments. Next, we introduce a score of merit for common attribute discovery based on the intuition that images perturbed along the same attribute will result in similar prediction changes, when measured through an “oracle” attribute classifier [23].

We first generate a batch of random samples from the latent spaces of client and reference GANs, and manipulate them along a common attribute direction $(\delta_n, \tilde{\delta}_n)$ inferred using xGA. In other words, we synthesize pairs of original and attribute-manipulated images from the two GANs and for each pair, we measure the discrepancy in the predictions from an “oracle” attribute classifier. Mathematically, this can be expressed as $a_i^n = |\mathcal{C}(\mathcal{G}_c(z_i, \delta_n)) - \mathcal{C}(\mathcal{G}_c(z_i))|$ and $\tilde{a}_j^n = |\mathcal{C}(\mathcal{G}_r(\tilde{z}_j, \tilde{\delta}_n)) - \mathcal{C}(\mathcal{G}_r(\tilde{z}_j))|$, where \mathcal{C} is the attribute classifier trained using the labeled CelebA dataset. Finally,

we define an alignment score that compares the expected prediction discrepancy across the two GANs using cosine similarity (higher value indicates alignment).

$$\mathcal{A}_{\text{score}} = \mathbb{E}_n \left[\cos \left(\mathbb{E}_i [a_i^n], \mathbb{E}_j [\tilde{a}_j^n] \right) \right], \quad (7)$$

where the inner expectations are w.r.t. the batch of samples and the outer expectation is w.r.t. the N_c common attributes.

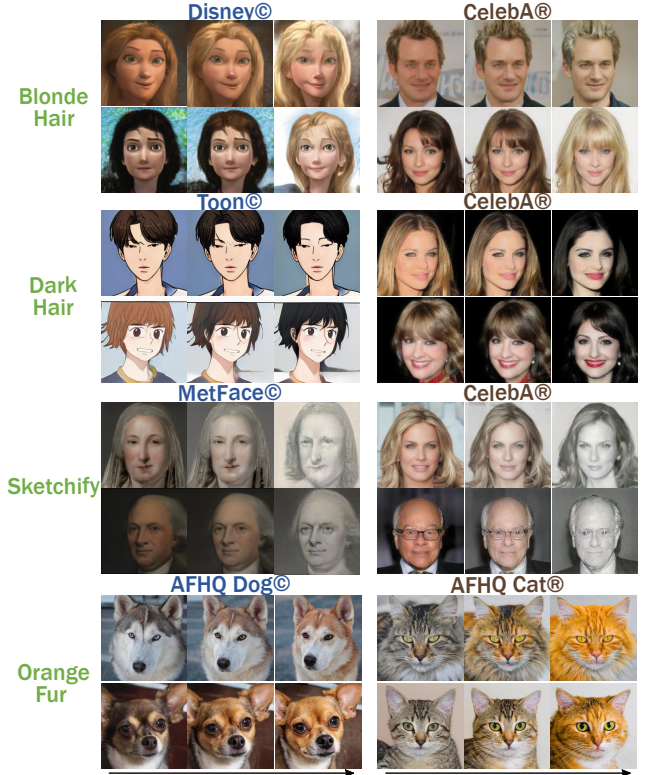


Figure 4. Visualizing common attributes discovered using xGA for different client-reference GAN pairs. For each case, we illustrate one common attribute (indicated by our description in green) with two random samples from the GAN latent space.

We implement 5 baseline approaches that apply state-of-the-art attribute discovery methods to the client and reference GANs (independently), and subsequently perform greedy, post-hoc alignment. In particular, we consider SeFa [34], Voynov [39], LatentCLR [43], Jacobian [40], and Hessian [28] methods for attribute discovery. Given the attributes for the two GANs, we use predictions from the “oracle” attribute classifier to measure the degree of alignment between every pair of directions. For example, the pair with the highest cosine similarity score is selected as the first common attribute. Next, we use the remaining latent directions to greedily pick the next attribute, and this process is repeated until we obtain $N_c = 12$ attributes. We compute the alignment score from (7) for all the methods and

Method	$\mathcal{A}_{\text{score}} (\uparrow)$	$\mathcal{R}_{\text{score}} (\uparrow)$
SeFa + G. S	0.382 ± 0.042	0.167 ± 0.165
Voynov + G. S	0.544 ± 0.033	0.254 ± 0.246
LatentCLR + G. S	0.543 ± 0.031	0.297 ± 0.326
Hessian + G. S	0.567 ± 0.065	0.224 ± 0.273
Jacobian + G. S	0.502 ± 0.024	0.233 ± 0.201
xGA	0.660 ± 0.147	0.411 ± 0.193

Table 1. **Common and Missing attribute discovery.** The average alignment scores from the 7 controlled CelebA experiments. Note, we report both the mean and standard deviations (\pm std) for each case, and “G. S” refers to the greedy strategy that we use for alignment.

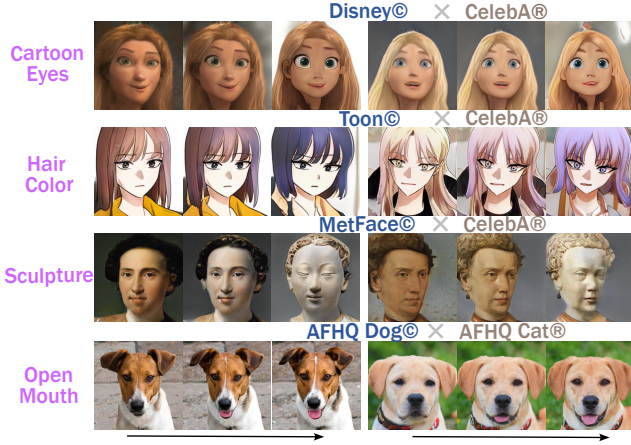


Figure 5. Visualizing novel attributes in different client GANs characterized by challenging distribution shifts with respect to the reference GAN (CelebA or AFHQ Cat-GANs). In each case, we show image manipulation in the attribute direction for two random sample from the latent space.

report results from the 7 controlled experiments in Table 1. Interestingly, we find that, despite using the “oracle” classifier for alignment, the performance of the baseline methods is significantly inferior to xGA. This clearly evidences the efficacy of our optimization strategy.

4.4. Evaluation: Novel/Missing Attribute Discovery

In this section, we study the effectiveness of xGA in discovering novel (only present in the client) and missing (only present in the reference model) attributes.

Qualitative results We first show results for novel attribute discovery for different client GANs in Figure 5. xGA produces highly intuitive results by identifying attributes that are unlikely to occur in the reference GAN. For example, “cartoon eyes” and “sculptures” are found to be unique to Disney and Met Faces GANs, when compared to CelebA. Next, we performed missing attribute discovery from the controlled CelebA experiments, where we know precisely which attribute is not encoded by the client GAN w.r.t the

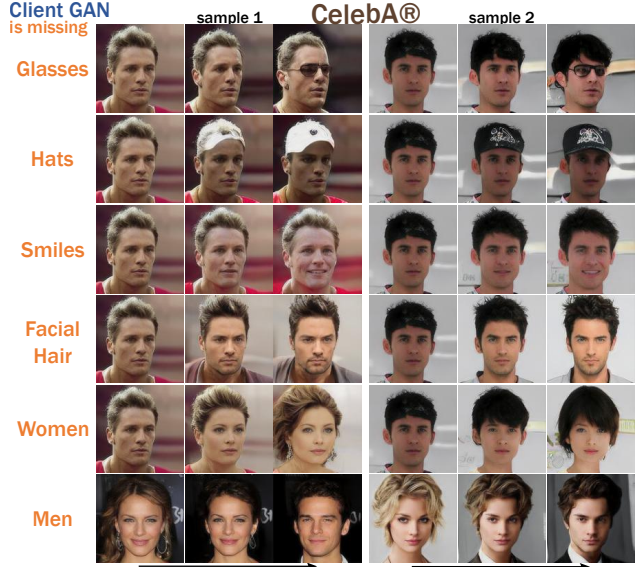


Figure 6. Using multiple clients trained with different subsets of CelebA data (one of the face attributes explicitly dropped), we find that, in all cases, xGA accurately recovers the missing attribute.

reference (standard CelebA StyleGANv2). As described earlier, the client models are always trained on a subset of data used by the reference model and by design, there are no novel attributes. Figure 6 shows examples for the different missing attributes. We find that xGA successfully reveals each of the missing client attributes, even though the data distributions $P_c(x)$ and $P_r(x)$ are highly similar (except for a specific missing attribute).

Quantitative results To benchmark xGA in missing attribute discovery, we use the 7 controlled CelebA client models and audit with respect to the reference CelebA GAN. We denote the set of attributes (one or more) which are explicitly excluded in each client model by \mathcal{M} . In order to evaluate how well xGA identifies the excluded attributes, we introduce a metric based on mean reciprocal rank (MRR) [30, 38]. For each of the N_m missing attributes from xGA, we compute the average semantic discrepancy from the “oracle” attribute classifier as,

$$a^n = \mathbb{E}_i[|\mathcal{C}(\mathcal{G}_c(z_i, \delta_n)) - \mathcal{C}(\mathcal{G}_c(z_i))|].$$

Denoting the rank of a missing attribute $m \in \mathcal{M}$ in the difference vector a^n as $\text{rank}(m, a^n)$, we can define the attribute recovery (for both missing/novelty) score as:

$$\mathcal{R}_{\text{Score}} = \mathbb{E}_m \left[\max_n \left(\frac{1}{\text{rank}(m, a^n)} \right) \right] \quad (8)$$

In Table 1, we show results for missing attribute discovery based on this score. We observe that xGA significantly outperforms all baselines in identifying the missing attribute across the suite of client GANs.

Method	$\mathcal{H}_{\text{score}} (\downarrow)$	$\mathcal{D}_{\text{score}} (\uparrow)$
SeFa [34]	4.006 ± 0.259	1.031 ± 0.077
LatentCLR [43]	2.348 ± 0.203	0.749 ± 0.929
Voynov [39]	2.508 ± 0.069	0.585 ± 0.725
Hessian [28]	2.707 ± 0.145	0.642 ± 0.795
Jacobian [40]	2.675 ± 0.070	0.661 ± 0.826
xGA (ViT)	1.988 ± 0.068	3.072 ± 3.845
xGA (MAE ViT)	2.102 ± 0.035	3.103 ± 3.814
xGA (CLIP ViT)	2.091 ± 0.041	3.135 ± 3.901
xGA (ResNet-50)	1.901 ± 0.060	3.111 ± 3.852
xGA (Clip ResNet-50)	2.033 ± 0.038	3.121 ± 3.863
xGA (advBN ResNet-50)	1.881 ± 0.057	3.153 ± 3.904

Table 2. **Choice of the feature space for attribute discovery.** Using an external feature space is superior to GAN’s native style space, in terms of both entropy ($\times 100$) and deviation metrics. In this experiment, we set $\mathcal{G}_r = \mathcal{G}_c$, and aggregate the metrics from the set of controlled CelebA StyleGANs.

4.5. Analysis

In this section, we examine the key components of xGA to understand its behavior better.

Impact of the Choice of \mathcal{F} We start by studying the choice of the external, feature space used to perform attribute discovery. For this analysis, we consider the case where we assume $\mathcal{G}_r = \mathcal{G}_c$, wherein xGA simplifies to the standard setting of attribute discovery with a single GAN model (set $\lambda_b = 0$), such as SeFa and latentCLR. We make an interesting observation that, using a robust latent space leads to improved diversity and disentanglement in the inferred attributes, when compared to the native latent space of StyleGAN. To quantify this behavior we consider two evaluation metrics based on the predictions for a batch of synthesized images $\mathcal{G}_c(z, \delta_n)$ from the “oracle” attribute classifier. First, for each latent direction δ_n , the average prediction entropy $\mathcal{H}_{\text{score}}$ [23] is defined as:

$$\mathcal{H}_{\text{score}} = \mathbb{E}_n \left[\mathbb{E}_i \left[\text{Entropy}(\mathcal{C}(\mathcal{G}_c(z_i, \delta_n))) \right] \right] \quad (9)$$

Second, the deviation in the predictions across all latent directions $\mathcal{D}_{\text{score}}$ is defined in (10), where K is the total number of attributes in the “oracle” classifier \mathcal{C} :

$$\mathcal{D}_{\text{score}} = \sum_{k=1}^K \text{Variance} \left[\left\{ \mathbb{E}_i [\mathcal{C}(\mathcal{G}_c(z_i, \delta_n))] \right\}_{n=1}^N \right]_k \quad (10)$$

When the entropy is low, it indicates that the semantic manipulation is concentrated to a specific attribute, and hence disentangled. On the other hand, when the deviation is high, it is reflective of the high diversity in the inferred latent directions.

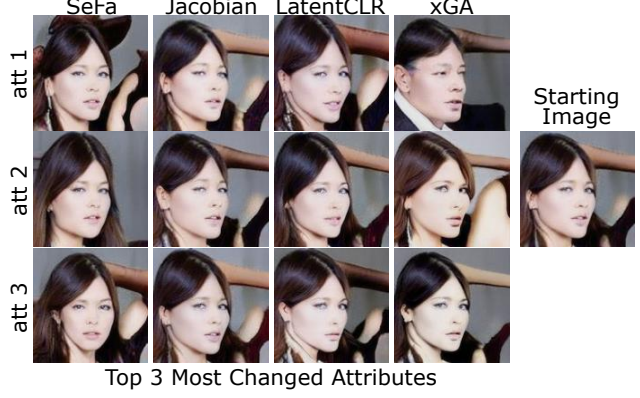


Figure 7. Comparing xGA on single GAN attribute discovery with existing approaches, we find that more diverse and novel attributes can be found simply by using an external feature space. We exploit this for effective alignment across two GAN models.

For this analysis, we considered the following feature extractors for implementing xGA: (i) vanilla ResNet-50 trained on ImageNet [15]; (ii) robust variant of ResNet-50 trained with advBN [35]; (iii) ResNet-50 trained via CLIP [31]. Table 2 shows the performance of the three feature extractors on attribute discovery with our 7 CelebA GANs trained using different data subsets. Note, we scale all entropy and diversity scores by 100 for ease of readability. We make a striking finding that, in terms both the entropy and deviation scores, performing attribute discovery in an external feature space is significantly superior to carrying out the optimization in the native style space (all baselines). As expected, LatentCLR produces the most disentangled attributes among the baselines, and regardless of the choice of \mathcal{F} , xGA leads to significant improvements. More importantly, the key benefit of xGA becomes more apparent from the improvements in the deviation score over the baselines. In the supplement, we include examples for the attributes inferred using all the methods. Finally, among the different choices for \mathcal{F} , the advBN ResNet-50 performs the best in terms of both metrics and hence it was used in all our experiments.

Single GAN Qualitative Results Figure 11 visualizes a shortened example of the top 3 attributes (induce most changes in the “oracle” classifier predictions). This example show a clear improvement by using a pretrained feature extractor, as xGA identifies the most diverse semantic changes. Complete results, all discovered attributes for all methods, are shown in the supplement.

Extending xGA to compare multiple GANs Though all our experiments used a client model w.r.t a reference, our method can be readily extended to perform comparative analysis of multiple GANs, with the only constraint arising from GPU memory since all generators need to be loaded into memory for optimization. We performed a




Figure 8. Common attributes identified using xGA with three different StyleGANs.

proof-of-concept experiment by discovering common attributes across 3 different independently trained StyleGANs as shown in Figure 8. For this setup, we expanded the cost function outlined in (3) to include 3 pairwise alignment terms from the 3 GANs to perform contrastive training, in addition to an extra independent term from the third model. While beyond scope for the current work, scaling xGA is an important direction for future work.

5. Discussion

We introduced the first cross-GAN auditing framework, which utilizes a novel optimization technique to jointly infer common, novel and missing attributes for a client GAN w.r.t any reference GAN. Through a large suite of datasets and GAN models, we demonstrate that the proposed method (i) consistently leads to higher quality (disentangled & diverse) attributes, (ii) effectively reveals shared attributes even across challenging distribution shifts, and (iii) accurately identifies the novel/missing attributes in our controlled experiments (i.e., known ground truth).

Limitations First, similar to other optimization-based attribute discovery approaches [39], [43], there is no guarantee that all prevalent factors are captured, though our controlled empirical studies clearly demonstrate the efficacy of xGA over existing approaches. Second, while using an external feature space enhances the performance of attribute discovery, this becomes an additional component that must be tuned. While we found advBN ResNet-50 to be a reasonable choice for a variety of face datasets (and AFHQ), a more systematic solution will expand the utility of our approach to other applications.

Acknowledgements

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. The project is directly supported by LDRD 22-ERD-006, and DOE HRRL. The manuscript is reviewed and released under LLNL-PROC-832985.

References

- [1] Amazon web services. <https://aws.amazon.com/>. Accessed: 2023-03-01. 2
- [2] Google cloud: Cloud computing services. <https://cloud.google.com/>. Accessed: 2023-03-01. 2
- [3] Microsoft azure: Cloud computing services. <https://azure.microsoft.com/en-us>. Accessed: 2023-03-01. 2
- [4] Ahmed Alaa, Boris Van Breugel, Evgeny S Saveliev, and Mihaela van der Schaar. How faithful is your synthetic data? sample-level metrics for evaluating and auditing generative models. In *International Conference on Machine Learning*, pages 290–306. PMLR, 2022. 2, 3
- [5] Jihye Back. Fine-tuning stylegan2 for cartoon face generation. *CoRR*, abs/2106.12445, 2021. 5
- [6] Viraj Bagal, Rishal Aggarwal, PK Vinod, and U Deva Priyakumar. Molgpt: Molecular generation using a transformer-decoder model. *Journal of Chemical Information and Modeling*, 2021. 1
- [7] David Bau, Jun-Yan Zhu, Jonas Wulff, William Peebles, Hendrik Strobelt, Bolei Zhou, and Antonio Torralba. Seeing what a gan cannot generate. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4502–4511, 2019. 2, 3
- [8] Brett K Beaulieu-Jones, Zhiwei Steven Wu, Chris Williams, Ran Lee, Sanjeev P Bhavnani, James Brian Byrd, and Casey S Greene. Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7):e005122, 2019. 1
- [9] Yumin Bian and Xiang-Qun Xie. Generative chemistry: drug discovery with deep learning generative models. *Journal of Molecular Modeling*, 27(3):1–18, 2021. 1
- [10] Jimmy S Chen, Aaron S Coyner, RV Paul Chan, M Elizabeth Hartnett, Darius M Moshfeghi, Leah A Owen, Jayashree Kalpathy-Cramer, Michael F Chiang, and J Peter Campbell. Deepfakes in ophthalmology: Applications and realism of synthetic retinal images from generative adversarial networks. *Ophthalmology Science*, 1(4):100079, 2021. 1
- [11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 3
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 1
- [13] Harshit Gupta, Thong H Phan, Jaejun Yoo, and Michael Unser. Multi-cryogan: Reconstruction of continuous conformations in cryo-em using generative adversarial networks. In *European Conference on Computer Vision*, pages 429–444. Springer, 2020. 1
- [14] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. GANspace: Discovering interpretable GAN controls. *arXiv preprint arXiv:2004.02546*, 2020. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-*

- ings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016. 8
- [16] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 2
- [17] Drew A Hudson and C. Lawrence Zitnick. Compositional transformers for scene generation. *Advances in Neural Information Processing Systems NeurIPS 2021*, 2021. 5, 12
- [18] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. 5, 12
- [19] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34, 2021. 1, 12
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 1, 2, 5
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of styleGAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 1
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [23] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 5, 6, 8
- [24] A. Menon and C. S. Ong. Linking losses for density ratio and class-probability estimation. In *Proceedings of the 33rd International Conference on Machine Learning*, page 304–313, 2016. 4, 5
- [25] Hyunha Nam and Masashi Sugiyama. Direct density ratio estimation with convolutional neural networks with application in outlier detection. *IEICE Transactions on Information and Systems*, E98.D(5):1073–1079, 2015. 4
- [26] Matthew Lyle Olson, Shusen Liu, Rushil Anirudh, Jayaraman J Thiagarajan, Weng-Keen Wong, and Peer-Timo Bremer. Unsupervised attribute alignment for characterizing distribution shift. In *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2021. 3
- [27] Matthew L Olson, Thuy-Vy Nguyen, Gaurav Dixit, Neale Ratzlaff, Weng-Keen Wong, and Minsuk Kahng. Contrastive identification of covariate shift in image data. In *2021 IEEE Visualization Conference (VIS)*, pages 36–40. IEEE, 2021. 3
- [28] William Peebles, John Peebles, Jun-Yan Zhu, Alexei Efros, and Antonio Torralba. The hessian penalty: A weak prior for unsupervised disentanglement. In *European Conference on Computer Vision*, pages 581–597. Springer, 2020. 2, 6, 8, 12
- [29] Justin NM Pinkney and Doron Adler. Resolution dependent gan interpolation for controllable image synthesis between domains. *arXiv preprint arXiv:2010.05334*, 2020. 3
- [30] Dragomir R Radev, Hong Qi, Harris Wu, and Weiguo Fan. Evaluating web-based question answering systems. In *LREC*. Citeseer, 2002. 7
- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 8
- [32] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. Closing the ai accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 33–44, 2020. 2, 3
- [33] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. InterfaceGAN: Interpreting the disentangled face representation learned by GANs. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 2
- [34] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in GANs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1532–1540, 2021. 2, 6, 8, 12
- [35] Manli Shu, Zuxuan Wu, Micah Goldblum, and Tom Goldstein. Encoding robustness to image style via adversarial feature perturbations. *Advances in Neural Information Processing Systems*, 34, 2021. 6, 8
- [36] Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density ratio estimation in machine learning*. Cambridge University Press, 2012. 4
- [37] Masashi Sugiyama, Taiji Suzuki, Shinichi Nakajima, Hisashi Kashima, Paul von Büna, and Motoaki Kawanabe. Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics*, 60(4):699–746, 2008. 4
- [38] Ellen M Voorhees et al. The trec-8 question answering track report. In *Trec*, volume 99, pages 77–82, 1999. 7
- [39] Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. In *International conference on machine learning*, pages 9786–9796. PMLR, 2020. 2, 6, 8, 9, 12
- [40] Yuxiang Wei, Yupeng Shi, Xiao Liu, Zhilong Ji, Yuan Gao, Zhongqin Wu, and Wangmeng Zuo. Orthogonal jacobian regularization for unsupervised disentanglement in image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6721–6730, 2021. 2, 6, 8, 13
- [41] Zongze Wu, Yotam Nitzan, Eli Shechtman, and Dani Lischinski. Stylealign: Analysis and applications of aligned stylegan models. *arXiv preprint arXiv:2110.11323*, 2021. 3
- [42] Tom Yan and Chicheng Zhang. Active fairness auditing. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine*

Learning Research, pages 24929–24962. PMLR, 17–23 Jul 2022. 3

- [43] Oğuz Kaan Yüksel, Enis Simsar, Ezgi Gülperi Er, and Pinar Yanardag. Latentclr: A contrastive learning approach for unsupervised discovery of interpretable directions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14263–14272, 2021. 2, 3, 4, 6, 8, 9, 12

A. Ablation study

First, we investigate the effect of the λ_b parameter on KLIEP loss that allows us to discover novel attributes. In addition to KLIEP loss presented in the main text, we analyze a model trained with simple log loss used to train binary classifiers to predict the likelihood of a given sample.

A.1. Log Loss Model

This model, we denote as LOG, is nearly identical to the DRE models except instead of a softplus final activation, it uses a sigmoid function $\sigma(x) = \frac{1}{1+e^{-x}}$. These models are used to classify whether a given feature belongs to $\mathcal{G}_c(z)$ or $\mathcal{G}_r(\bar{z})$. We pre-train two separate LOG models to approximate $\hat{\gamma}_c(x) = \hat{p}_c(x)$, and $\hat{\gamma}_r(x) = \hat{p}_r(x)$, where we treat \mathcal{G}_c as $P(x|Y=1)$ and \mathcal{G}_r as $P(x|Y=0)$. These LOG models are learned using simple 2-layer MLPs, $f_{LOG}^c(\cdot)$, $f_{LOG}^r(\cdot)$, such that

$$\hat{\gamma}_c(z) = f_{LOG}^c(\mathcal{F}(\mathcal{G}_c(z))) \text{ and } \hat{\gamma}_r(\bar{z}) = f_{LOG}^r(\mathcal{F}(\mathcal{G}_r(\bar{z}))), \quad (11)$$

where \mathcal{F} is the same Encoder model used in main paper’s equation 3.

The loss used for training the LOG models is defined as follows:

$$\mathcal{L}_{Log}^c = \frac{1}{T_2} \sum_{j=1}^{T_2} -\log(1 - \hat{\gamma}_c(\bar{z}_j)) + \frac{1}{T_1} \sum_{i=1}^{T_1} -\log(\hat{\gamma}_c(z_i)) \quad (12)$$

where \bar{z}_j and z_i are random samples drawn from the latent space of each generator. The loss term for the second model LOG model is

$$\mathcal{L}_{Log}^r = \frac{1}{T_1} \sum_{j=1}^{T_1} -\log(1 - \hat{\gamma}_r(\bar{z}_j)) + \frac{1}{T_2} \sum_{i=1}^{T_2} -\log(\hat{\gamma}_r(z_i)) \quad (13)$$

The LOG models $f_{LOG}^1(\cdot)$, $f_{LOG}^2(\cdot)$ are trained to minimize \mathcal{L}_{Log}^1 , \mathcal{L}_{Log}^2 respectively.

Finally, the trained LOG models are used to minimize the loss in equation 7 (rather than DRE models); the objective in equation 7 remains the same.

A.2. Missing attribute ablation study results

Table 3 illustrates the missing attribute discovery score for each CelebA split versus full CelebA. With $\lambda = 0$ (i.e. ignoring the DRE loss), the attribute discovery process has difficulty capturing some missing attributes. When using a regularization model trained with Log-loss, the results are consistently worse than DRE, sometimes even worse than with $\lambda = 0$. The KLIEP loss model, on the other hand, performs consistently better for all lambda values > 0 .

λ	$\mathcal{R}_{\text{Score}}$ (DRE loss) (\uparrow)	$\mathcal{R}_{\text{Score}}$ (Log loss) (\uparrow)
0	0.42 ± 0.38	0.42 ± 0.38
0.1	0.61 ± 0.35	0.37 ± 0.41
0.2	0.54 ± 0.33	0.44 ± 0.39
0.5	0.57 ± 0.40	0.45 ± 0.38
1	0.61 ± 0.33	0.40 ± 0.40
5	0.57 ± 0.39	0.34 ± 0.32

Table 3. The effect on the unique direction score when modifying the regularization λ on the average $\mathcal{R}_{\text{Score}}$ (\pm std) for the the 7 CelebA pairwise leave-attribute-out experiments using a Robust ResNet-50 encoder.




Figure 9. An example of applying our method to two generative models trained on the same dataset (FFHQ). We find ProgGAN and GANformer are able to find some alignment, and that the newer model (GANformer) is better at capturing the full data distribution of FFHQ (Missing) whereas ProgGAN is prone to generating non-realistic images (Novel).

B. Same dataset, different architecture

To verify the effectiveness of xGA at comparing models trained on the same dataset with different configurations, we perform two sets of experiments. We use Prog-GAN [18] (client) and GANformer [17] (reference) trained on the FFHQ dataset. Figure 9 shows an example of how these two GANs can be aligned, and how the novel/missing attribute reflects each GAN’s capacity to learn the data distribution. We also use two configurations of a StyleGAN3 [19] trained on FFHQ. Figure 10 shows how translation equivariance is preserved in both models, whereas only the StyleGAN3-r model is rotationally equivariant.

C. Single GAN results

Here we present the full training of all learned directions for each of our methods using the same starting point from CelebA GAN. Figure 11 visualizes a shortened example of the top 3 attributes (induce most changes in the “oracle” classifier predictions) and it is clear that xGA identifies the most diverse semantic changes. Complete results can be seen as follows:




Figure 10. A few examples from our experiment applying xGA between two StyleGAN3 models, both trained on FFHQ, but with different model configurations. As expected both models having translation equivariance, and the rotation equivariance is missing from the translation model.



Figure 11. Comparing xGA on single GAN attribute discovery with existing approaches, we find that more diverse and novel attributes can be found simply by using an external feature space. We exploit this for effective alignment across two GAN models. Complete examples for all methods are provided below.

1. Sefa [34]: Figure 22
2. LatentCLR [43]: Figure 20
3. Voynov [39]: Figure 21
4. Hessian [28]: Figure 18

5. Jacobian [40]: Figure 19
6. xGA (ImageNet ResNet-50): figure 13
7. xGA (advBN ResNet-50): figure 14
8. xGA (CLIP ResNet-50): figure 15

We visualize both positive and negative directions for every model. Even though xGA and LatentCLR are not directly trained for negative directions, we find these attributes to be semantically meaningful and interesting.

Next we present an example where we compare two LatentCLR models trained on different GANs where the reference GAN is CelebA and client GAN is CelebA without Hats. We sort all the directions by most similar (as described in the main paper) and show an example of the results in Figure 12, finding many similarities, but no dedicated Hat attribute in the reference GAN. Showing how without the dedicated constraint of the DRE models, finding missing attributes is difficult.

D. Expanded Qualitative Results

Here we present many additional examples of shared directions between two GANs, and novel/missing directions from a few different GAN pairs that contain subset of the CelebA dataset. We introduce a new GAN (anime), as it produces interesting common, missing, and novel attributes, though the GAN itself produces lower quality images than other models, and as such we leave it here in the supplement. The figures are arranged as follows:

1. Common attributes: CelebA (reference) and Metface (client) sketch (23), formal (24), and curly hair (25)
2. Common attributes: Anime (client) and Toon (reference) purple hair (26), orange/brown hair (27), open mouths (28), and smiling (29); missing attributes of green hair / lipstick (30)
3. Common attributes: CelebA (reference) and Disney (client) blonde hair (31), and brown hair (32); novel Disney attributes of turning green / cartoonish eyes (33)
4. Additional missing attributes from different CelebA client GANs, with CelebA reference GAN (34)

E. Expanded Quantitative Results

First we present the results for using ViT-based feature extractors in table 6. We include 3 different pretrained models: one original trained on ImageNet, CLIP, and MAE. While ViT does well for entropy metric, it performs poorly for cross model based experiments.

Next, we present the entire results for our missing attribute quantitative experiments. To recap these experiments, we use the 7 controlled CelebA models which are missing one or more attributes (hat, glasses, male, female, beard, beards—hats, and smiles—glasses—ties) and treat them as the client model; we audit these models with respect to the reference CelebA GAN. The 7 missing attribute experiments are shown in table 4, where we can see xGA performs well (e.g., easily finding the missing glasses attribute). The 7 attribute alignment experiments are shown in table 5, where again we see xGA with a robust resnet performs well, especially when the client GAN is missing multiple attributes (e.g., client GAN is missing beards and hats).

For completion’s sake, we run pairwise experiments between each GAN, treating each GAN as reference versus the other 7 GANs, which results in a total of 56 client/reference paired experiments. We report these comprehensive results in the following tables (where rows are reference GAN and columns are the client): 11, 10, 9, 8, 12, 13, 14, 15, 16, 17, and 18. We also compute the the common attribute results experiments in the following tables: 22,21,20,23, 24, 25, 26, 27, 28, 29, and 30.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
SeFa	0.143	0.143	0.045	0.111	0.063	0.278	0.189
jacobian	0.478	0.536	0.086	0.390	0.388	0.120	0.287
Hessian	1.000	1.000	0.056	0.167	0.167	0.096	0.407
LatentCLR	1.000	1.000	0.250	0.333	0.200	0.153	0.537
Voynov	0.500	1.000	0.333	0.050	0.500	0.153	0.259
xGA (ResNet-50)	1.000	1.000	0.333	0.500	0.200	0.167	0.465
xGA (Clip ResNet-50)	1.000	1.000	1.000	0.200	0.200	0.108	0.383
xGA (advBN ResNet-50)	1.000	1.000	0.250	1.000	0.063	0.183	0.401

Table 4. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$), where CelebA GAN is the reference.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
SeFa	0.413	0.458	0.372	0.374	0.314	0.387	0.355
Hessian	0.475	0.489	0.652	0.598	0.618	0.615	0.525
LatentCLR	0.519	0.511	0.556	0.533	0.512	0.593	0.579
Voynov	0.566	0.477	0.567	0.555	0.570	0.562	0.513
Jacobian	0.523	0.452	0.505	0.528	0.519	0.491	0.495
xGA (ResNet-50)	0.457	0.403	0.740	0.792	0.461	0.643	0.489
xGA (Clip ResNet-50)	0.753	0.451	0.772	0.791	0.894	0.580	0.656
xGA (advBN ResNet-50)	0.615	0.357	0.750	0.825	0.619	0.803	0.649

Table 5. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$), where CelebA GAN is the reference

Method / Model	$\mathcal{H}_{\text{score}}$ (\downarrow)	$\mathcal{A}_{\text{score}}$ (\uparrow)	$\mathcal{R}_{\text{score}}$ (\uparrow)
xGA + ViT	1.988 ± 0.068	0.377 ± 0.090	0.249 ± 0.217
xGA + ViT + MAE	2.102 ± 0.035	0.349 ± 0.089	0.194 ± 0.197
xGA + ViT + Clip	2.091 ± 0.041	0.397 ± 0.122	0.268 ± 0.195

Table 6. ViT-based extractors results. The average entropy scores for all 8 CelebA experiments, the average alignment scores ($\mathcal{A}_{\text{score}}$) for the CelebA pairwise experiments, and the average recovery scores ($\mathcal{R}_{\text{score}}$) for the CelebA pairwise leave-attribute-out experiments (\pm std)

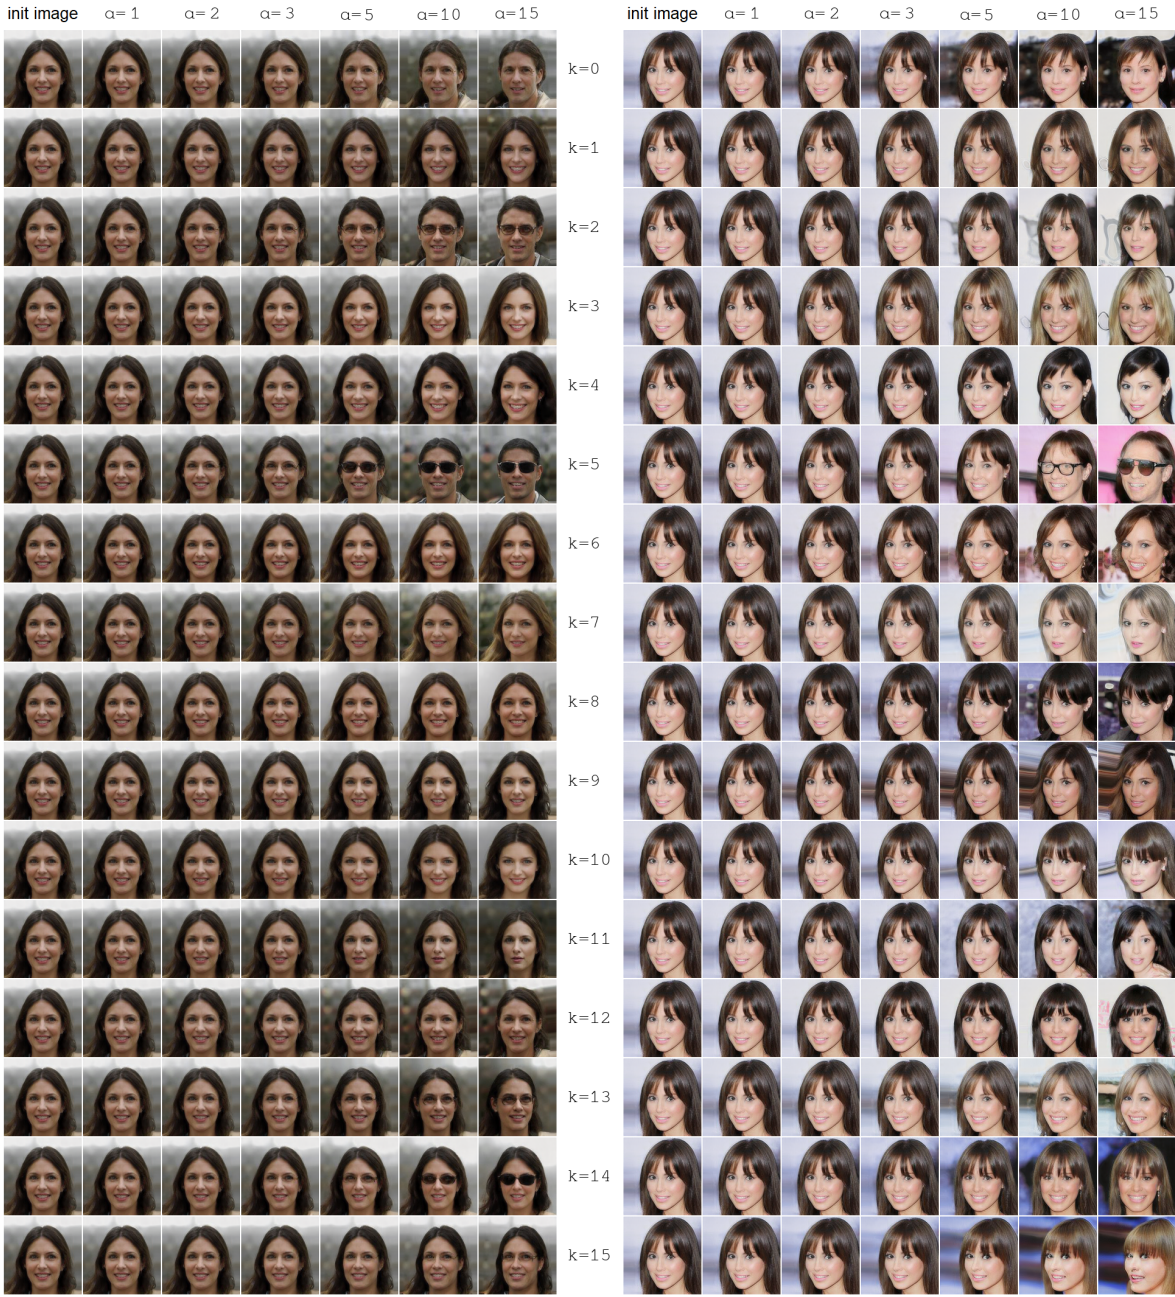


Figure 12. All 16 directions of two single GAN LatentCLR models trained on different GANs where the reference GAN is CelebA and client GAN is CelebA without Hats. The attributes are sorted by most similar ($k = 0$) to least similar ($k = 15$). While there are some major similarities (short hair $k = 0$, eyeglasses $k = 5$), the lack of a dedicated constraint for finding hats shows the flaws with this approach. This example is a clear demonstration that without a dedicated cross-model constraint (i.e., using the DRE models), finding missing attributes is difficult.




Figure 13. All 16 learned attributes of a Vanilla ResNet model.

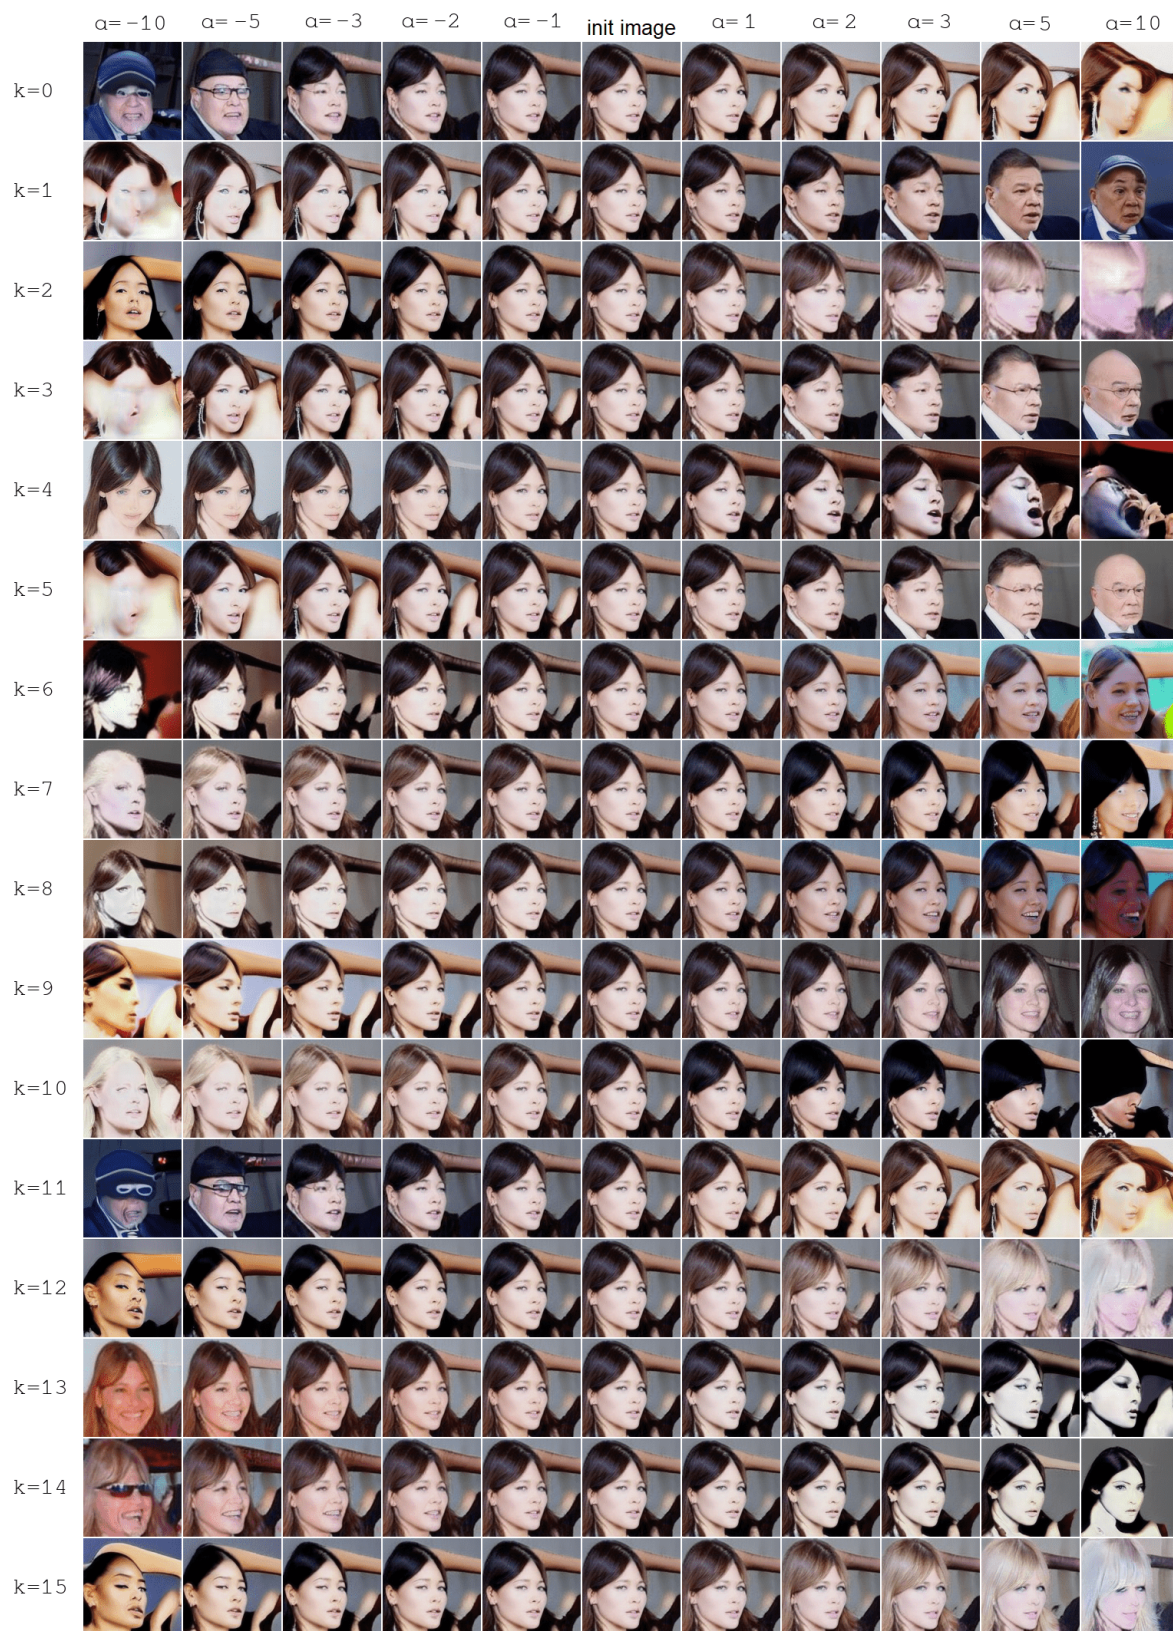


Figure 14. All 16 learned attributes of a robust ResNet model.



Figure 15. All 16 learned attributes of a clip ResNet model.

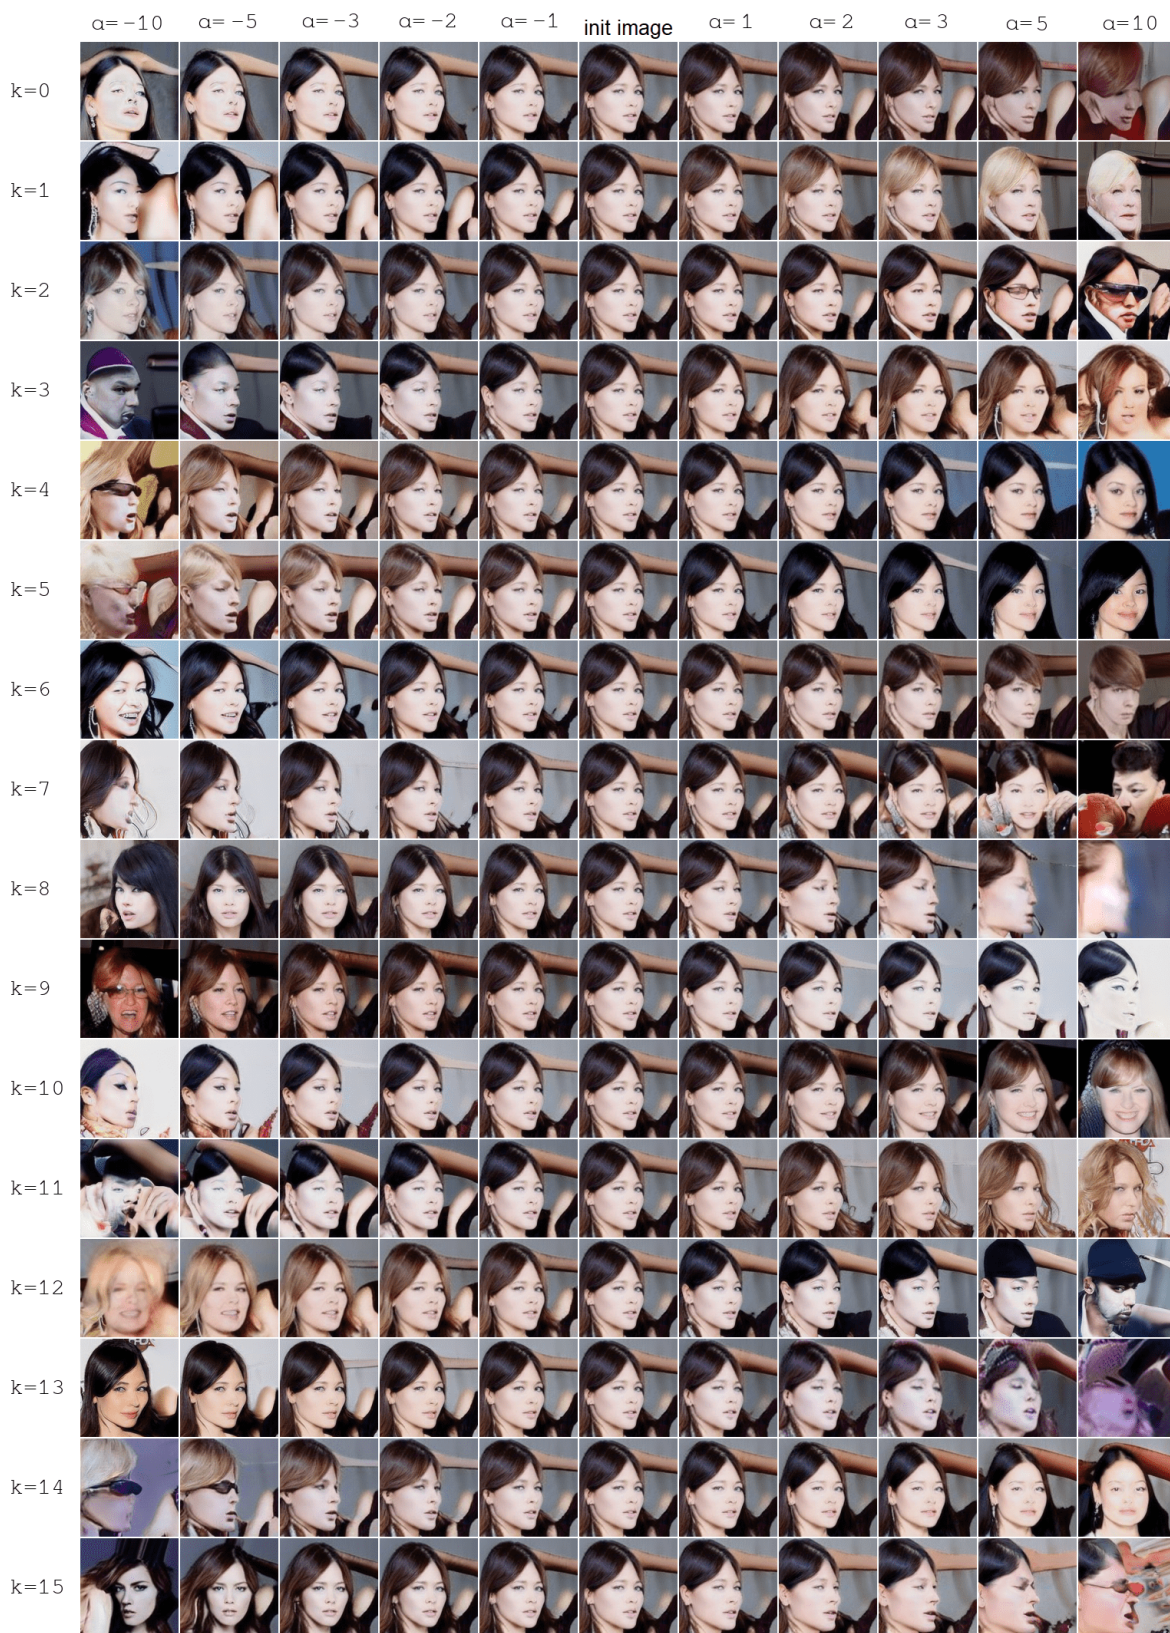


Figure 16. All 16 learned attributes of a Attribute Classifier ResNet model.




Figure 17. All 16 learned attributes of the original LatentCLR model (using global directions, rather than conditional).



Figure 18. All 16 learned attributes of the Hessian method.




Figure 19. All 16 learned attributes of the Jacobian method.



Figure 20. All 16 learned attributes of the original LatentCLR model with conditional directions.




Figure 21. All 16 learned attributes of the Voynov method.



Figure 22. The top 16 learned attributes of the SeFa method.




Figure 23. Examples of common sketch attribute between CelebA (Top) and Metfaces (Bottom).



Figure 24. Examples of common formal-wear attribute between CelebA (Top) and Metfaces (Bottom).



Figure 25. Examples of common white, curly hair attribute between CelebA (Top) and Metfaces (Bottom).




Figure 26. Examples of common purple hair attribute between Anime (Top) and Toon (Bottom).




Figure 27. Examples of common orange/brown hair attribute between Anime (Top) and Toon (Bottom).




Figure 28. Examples of common open-mouth attribute between Anime (Top) and Toon (Bottom).




Figure 29. Examples of common smiling attribute between Anime (Top) and Toon (Bottom).




Figure 30. Examples of novel green hair attribute from Anime (Top) and the missing lipstick attribute from Toon (Bottom).



Figure 31. Examples of common blonde attribute between CelebA (Top) and Disney (Bottom).

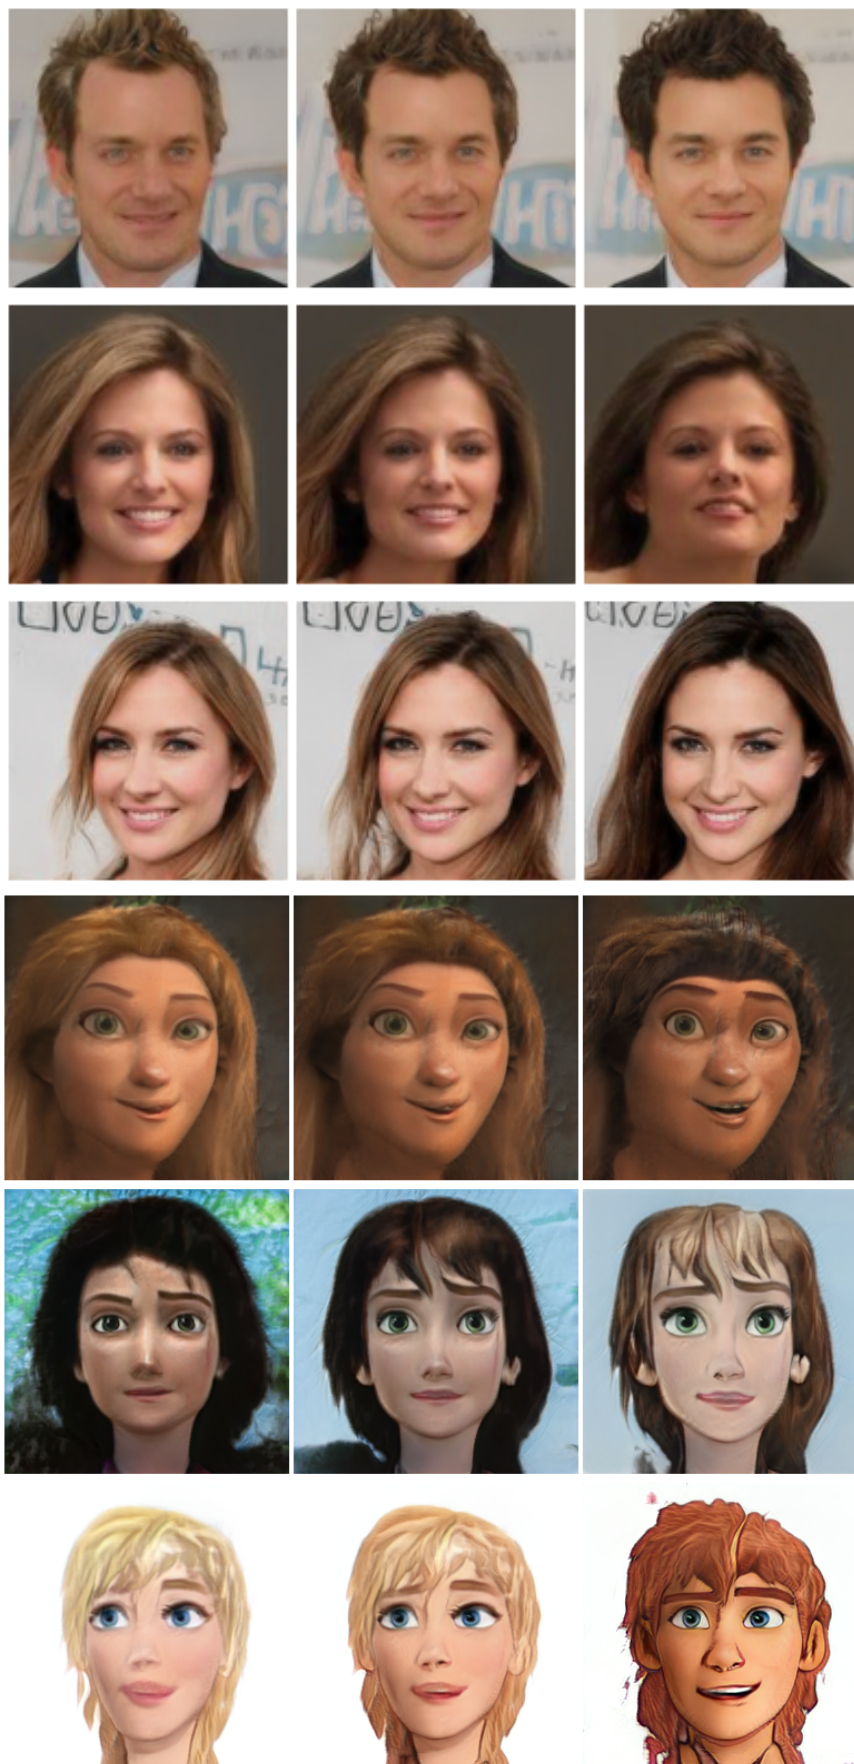


Figure 32. Examples of common brown hair attribute between CelebA (Top) and Disney (Bottom).

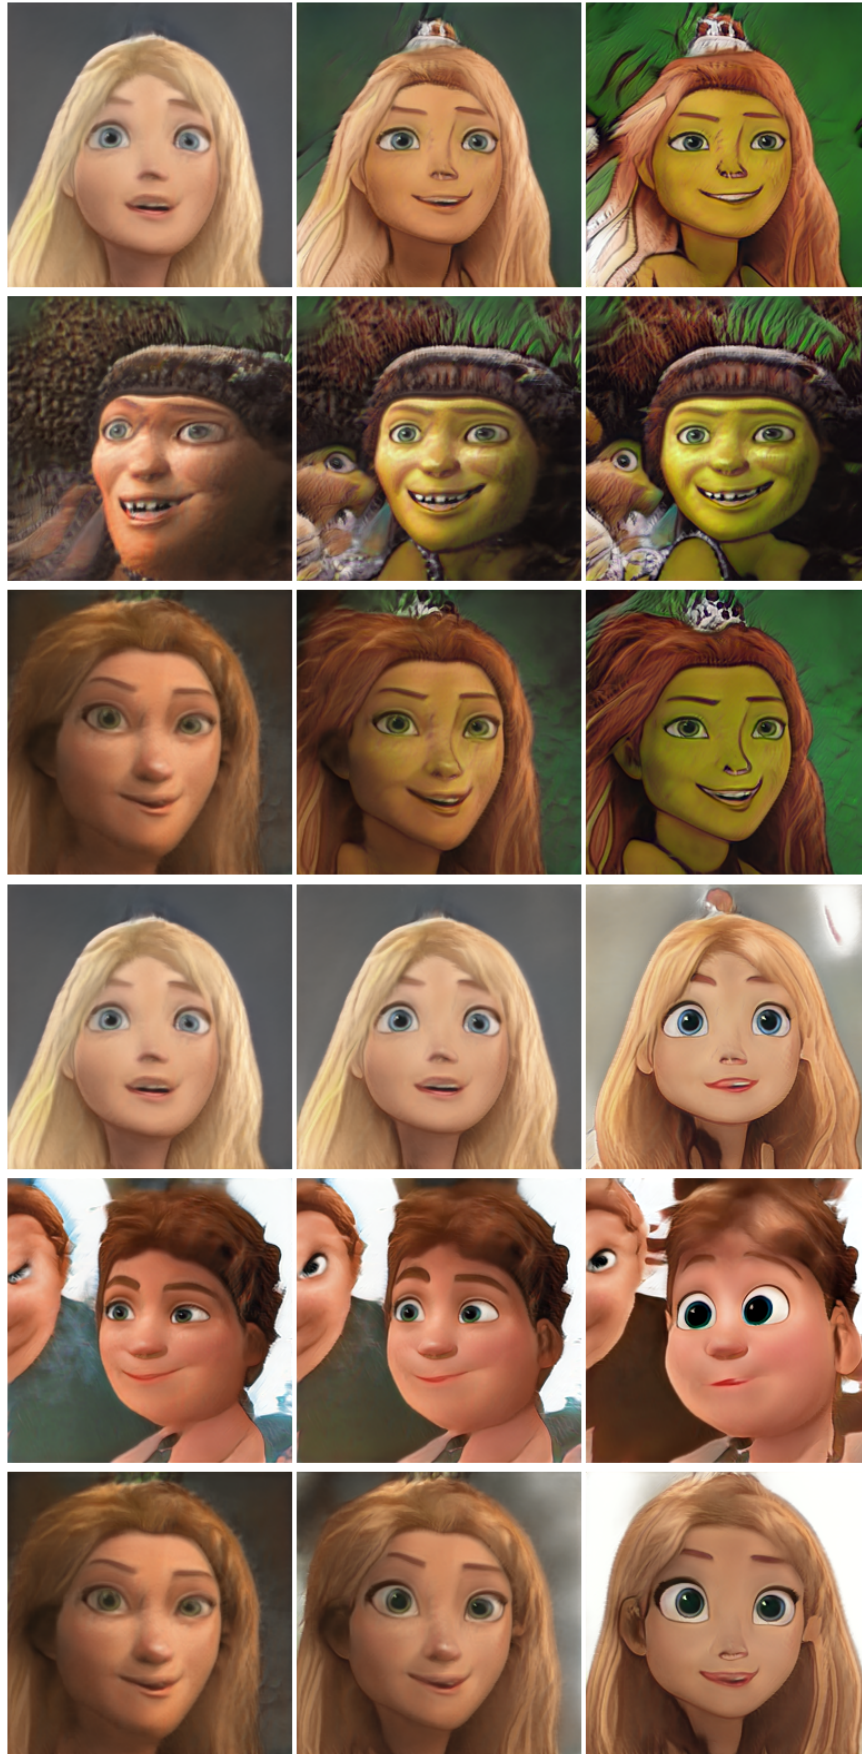


Figure 33. Two examples of novel Disney attributes: making princesses ogre-like and large cartoonish eyes.



Figure 34. Examples of various missing attributes from full CelebA GAN against three different attribute splits: **(Left 3)** the missing eyeglass attribute, **(Middle 3)** the missing hats attribute and **(Right 3)** the missing eyeglass, smiling, and attempting to identify the necktie attribute.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA		0.143	0.143	0.045	0.111	0.063	0.278	0.189
Female	0.000		0.167	1.000	0.167	0.200	0.170	0.417
Male	0.000	0.059		0.250	0.111	0.250	0.188	0.303
No Hats	0.000	0.056	0.111		0.083	0.067	0.188	0.267
No Glasses	0.000	0.059	0.125	0.333		0.042	0.306	0.203
No Beards	0.000	0.333	0.083	0.043	0.063		0.185	0.107
No Beard No Hats	0.000	0.143	0.100	0.125	0.143	0.125		0.511
No Glasses No Smiles No Ties	0.000	0.053	0.333	0.250	0.500	0.059	0.096	

Table 7. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the SeFa method.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	0.478	0.536	0.086	0.390	0.388	0.120	0.287
Female	0	-	0.246	0.048	0.050	0.046	0.048	0.279
Male	0	0.586	-	0.124	0.333	0.733	0.384	0.405
No Hats	0	0.251	0.240	-	0.097	0.180	0.108	0.313
No Glasses	0	0.585	0.542	0.048	-	0.114	0.100	0.218
No Beards	0	0.290	0.583	0.069	0.080	-	0.066	0.271
No Beard No Hats	0	0.373	0.396	0.034	0.189	0.049	-	0.297
No Glasses No Smiles No Ties	0	0.448	0.667	0.055	0.033	0.537	0.269	-

Table 8. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the Jacobian loss.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	1.000	1.000	0.056	0.167	0.167	0.096	0.407
Female	0	-	0.200	0.056	0.045	0.036	0.057	0.364
Male	0	0.333	-	0.077	1.000	1.000	0.300	0.300
No Hats	0	0.500	0.250	-	0.083	0.071	0.042	0.370
No Glasses	0	0.333	0.250	0.083	-	0.111	0.071	0.150
No Beards	0	0.125	0.167	0.071	0.063	-	0.046	0.218
No Beard No Hats	0	0.167	0.333	0.032	0.071	0.053	-	0.375
No Glasses No Smiles No Ties	0	1.000	0.333	0.200	0.048	0.143	0.102	-

Table 9. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the Hessian loss.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	1.000	1.000	0.250	0.333	0.200	0.153	0.537
Female	0	-	0.250	0.048	0.034	0.050	0.044	0.137
Male	0	1.000	-	0.143	0.500	0.333	0.267	0.242
No Hats	0	0.250	1.000	-	0.125	0.167	0.077	0.158
No Glasses	0	0.333	1.000	0.038	-	0.250	0.525	0.153
No Beards	0	0.125	0.500	0.045	0.063	-	0.049	0.381
No Beard No Hats	0	1.000	1.000	0.043	0.167	0.091	-	0.389
No Glasses No Smiles No Ties	0	1.000	0.500	0.067	0.033	0.125	0.072	-

Table 10. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the LatentCLR loss.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	0.500	1.000	0.333	0.050	0.500	0.153	0.259
Female	0	-	0.143	0.045	0.167	0.038	0.061	0.410
Male	0	0.500	-	0.250	0.333	0.333	0.306	0.256
No Hats	0	0.500	0.143	-	0.091	0.500	0.139	0.511
No Glasses	0	0.333	1.000	0.050	-	0.167	0.094	0.377
No Beards	0	0.167	0.250	0.167	0.091	-	0.052	0.370
No Beard No Hats	0	0.500	0.500	0.034	0.050	0.034	-	0.386
No Glasses No Smiles No Ties	0	0.143	0.500	0.143	0.029	1.000	0.286	-

Table 11. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the voynov loss.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	1.000	1.000	0.333	0.500	0.200	0.167	0.465
Female	0	-	0.333	0.053	0.067	0.034	0.046	0.381
Male	0	1	-	0.059	0.250	0.083	0.082	0.521
No Hats	0	1	0.5	-	0.143	0.071	0.062	0.460
No Glasses	0	1	0.5	0.059	-	0.091	0.081	0.377
No Beards	0	1	1	1	0.091	-	0.154	0.378
No Beard No Hats	0	1	1	0.033	0.05	0.042	-	0.382
No Glasses No Smiles No Ties	0	1	1	0.083	0.029	0.333	0.122	-

Table 12. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the ImageNet ResNet.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	1.000	0.333	1.000	0.091	1.000	0.525	0.390
Female	0	-	0.200	0.042	0.059	0.038	0.036	0.198
Male	0	1	-	0.063	0.143	0.056	0.306	0.492
No Hats	0	1	0.333	-	0.333	1.000	0.516	0.365
No Glasses	0	1	0.333	0.1	-	0.500	0.270	0.357
No Beards	0	1	1	0.053	0.111	-	0.042	0.211
No Beard No Hats	0	1	1	0.033	0.071	0.048	-	0.212
No Glasses No Smiles No Ties	0	1	1	1	0.03	0.5	0.274	-

Table 13. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the Attribute Classifier ResNet.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	1.000	1.000	0.250	1.000	0.063	0.183	0.401
Female	0	-	0.333	0.500	0.042	0.048	0.047	0.363
Male	0	0.5	-	0.077	0.500	0.111	0.563	0.419
No Hats	0	1	1	-	1.000	0.143	0.517	0.423
No Glasses	0	1	0.333	0.034	-	0.333	0.200	0.361
No Beards	0	1	1	0.143	0.045	-	0.113	0.370
No Beard No Hats	0	1	1	0.028	0.5	0.038	-	0.376
No Glasses No Smiles No Ties	0	1	1	0.056	0.033	1	0.556	-

Table 14. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the Robust ResNet.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	1.000	1.000	1.000	0.200	0.200	0.108	0.383
Female	0	-	0.250	0.028	0.200	0.063	0.131	0.194
Male	0	0.143	-	0.028	0.200	0.063	0.131	0.055
No Hats	0	1	1	-	0.500	0.125	0.133	0.231
No Glasses	0	1	1	0.037	-	0.100	0.148	0.397
No Beards	0	1	1	0.083	0.333	-	0.098	0.373
No Beard No Hats	0	1	1	0.1	0.2	0.029	-	0.371
No Glasses No Smiles No Ties	0	1	1	0.063	0.125	0.167	0.205	-

Table 15. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the CLIP ResNet.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	1.000	1.000	0.200	0.091	0.050	0.276	0.363
Female	0	-	0.333	0.040	0.045	0.056	0.054	0.140
Male	0	1	-	0.500	0.111	0.063	0.140	0.199
No Hats	0	0.5	1	-	0.200	0.040	0.076	0.369
No Glasses	0	0.333	1	0.111	-	0.063	0.042	0.388
No Beards	0	0.143	0.143	0.125	0.1	-	0.086	0.209
No Beard No Hats	0	0.143	0.5	0.031	0.053	0.032	-	0.192
No Glasses No Smiles No Ties	0	1	1	0.033	0.125	0.125	0.108	-

Table 16. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the ImageNet ViT.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	0.500	1.000	0.027	1.000	0.067	0.264	0.387
Female	0	-	0.200	0.034	0.032	0.036	0.048	0.188
Male	0	0.333	-	0.091	0.125	0.045	0.229	0.232
No Hats	0	1	1	-	0.063	0.067	0.170	0.365
No Glasses	0	1	0.5	0.067	-	0.053	0.066	0.189
No Beards	0	0.5	1	0.091	0.333	-	0.073	0.127
No Beard No Hats	0	0.2	0.5	0.029	0.056	0.038	-	0.369
No Glasses No Smiles No Ties	0	1	1	0.032	0.125	0.037	0.107	-

Table 17. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the CLIP ViT.

	Full CelebA	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	-	0.200	0.250	0.100	0.077	0.077	0.156	0.068
Female	0	-	0.200	0.038	0.056	0.033	0.044	0.119
Male	0	1	-	0.053	0.143	0.042	0.104	0.511
No Hats	0	0.091	1	-	0.063	0.028	0.046	0.194
No Glasses	0	0.25	1	0.067	-	0.034	0.047	0.081
No Beards	0	0.5	1	0.045	0.083	-	0.072	0.356
No Beard No Hats	0	0.25	0.333	0.067	0.059	0.034	-	0.068
No Glasses No Smiles No Ties	0	1	0.5	0.027	0.091	0.043	0.163	-

Table 18. The full results for the recovery scores ($\mathcal{R}_{\text{score}}$) of the MAE ViT.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.413	0.458	0.3715	0.374	0.314	0.387	0.355
Female	-	0.329	0.3615	0.320	0.352	0.360	0.334
Male	-	-	0.3732	0.426	0.364	0.352	0.381
No Hats	-	-	-	0.400	0.300	0.325	0.379
No Glasses	-	-	-	-	0.336	0.305	0.343
No Beards	-	-	-	-	-	0.302	0.295
No Beard No Hats	-	-	-	-	-	-	0.341

Table 19. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the SeFa method.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.475	0.489	0.652	0.598	0.618	0.615	0.525
Female	-	0.367	0.523	0.457	0.559	0.559	0.345
Male	-	-	0.466	0.401	0.372	0.431	0.450
No Hats	-	-	-	0.512	0.620	0.505	0.518
No Glasses	-	-	-	-	0.556	0.502	0.503
No Beards	-	-	-	-	-	0.563	0.477
No Beard No Hats	-	-	-	-	-	-	0.439

Table 20. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the Hessian loss.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.519	0.511	0.556	0.533	0.512	0.593	0.579
Female	-	0.412	0.452	0.513	0.550	0.613	0.388
Male	-	-	0.373	0.474	0.425	0.448	0.426
No Hats	-	-	-	0.460	0.485	0.576	0.482
No Glasses	-	-	-	-	0.484	0.630	0.491
No Beards	-	-	-	-	-	0.550	0.501
No Beard No Hats	-	-	-	-	-	-	0.496

Table 21. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the LatentCLR loss.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.566	0.477	0.567	0.555	0.570	0.562	0.513
Female	-	0.430	0.606	0.609	0.577	0.592	0.505
Male	-	-	0.508	0.503	0.429	0.456	0.453
No Hats	-	-	-	0.596	0.553	0.583	0.553
No Glasses	-	-	-	-	0.572	0.559	0.564
No Beards	-	-	-	-	-	0.616	0.567
No Beard No Hats	-	-	-	-	-	-	0.499

Table 22. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the Voynov method.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.523	0.452	0.505	0.528	0.519	0.491	0.495
Female	-	0.413	0.481	0.522	0.490	0.511	0.449
Male	-	-	0.488	0.451	0.443	0.409	0.384
No Hats	-	-	-	0.498	0.515	0.523	0.465
No Glasses	-	-	-	-	0.497	0.517	0.503
No Beards	-	-	-	-	-	0.544	0.484
No Beard No Hats	-	-	-	-	-	-	0.512

Table 23. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the Jacobian loss.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.457	0.403	0.740	0.792	0.461	0.643	0.489
Female	-	0.409	0.651	0.720	0.599	0.483	0.444
Male	-	-	0.508	0.377	0.328	0.360	0.390
No Hats	-	-	-	0.611	0.778	0.414	0.560
No Glasses	-	-	-	-	0.698	0.659	0.649
No Beards	-	-	-	-	-	0.652	0.556
No Beard No Hats	-	-	-	-	-	-	0.492

Table 24. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the ImageNet Trained ResNet.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.069	0.272	0.419	0.494	0.556	0.543	0.014
Female	-	0.155	0.411	0.230	0.191	0.354	0.056
Male	-	-	0.248	0.392	0.086	0.338	0.231
No Hats	-	-	-	0.370	0.279	0.346	0.569
No Glasses	-	-	-	-	0.494	0.405	0.447
No Beards	-	-	-	-	-	0.331	0.510
No Beard No Hats	-	-	-	-	-	-	0.403

Table 25. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the Attribute Classifier ResNet.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.615	0.357	0.750	0.825	0.619	0.803	0.649
Female	-	0.439	0.643	0.644	0.404	0.565	0.525
Male	-	-	0.384	0.444	0.417	0.175	0.289
No Hats	-	-	-	0.508	0.310	0.744	0.641
No Glasses	-	-	-	-	0.717	0.682	0.557
No Beards	-	-	-	-	-	0.568	0.495
No Beard No Hats	-	-	-	-	-	-	0.474

Table 26. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the Robust ResNet.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.753	0.451	0.772	0.791	0.894	0.580	0.656
Female	-	0.283	0.733	0.684	0.639	0.474	0.337
Male	-	-	0.371	0.435	0.396	0.337	0.314
No Hats	-	-	-	0.815	0.679	0.715	0.505
No Glasses	-	-	-	-	0.748	0.608	0.623
No Beards	-	-	-	-	-	0.706	0.507
No Beard No Hats	-	-	-	-	-	-	0.723

Table 27. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the CLIP ResNet.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.308	0.417	0.433	0.486	0.455	0.409	0.416
Female	-	0.251	0.348	0.476	0.468	0.495	0.365
Male	-	-	0.168	0.275	0.304	0.331	0.363
No Hats	-	-	-	0.409	0.484	0.499	0.390
No Glasses	-	-	-	-	0.347	0.363	0.417
No Beards	-	-	-	-	-	0.344	0.164
No Beard No Hats	-	-	-	-	-	-	0.363

Table 28. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the ImageNet ViT.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.418	0.373	0.496	0.358	0.579	0.539	0.499
Female	-	0.169	0.343	0.399	0.368	0.504	0.258
Male	-	-	0.183	0.381	0.213	0.187	0.225
No Hats	-	-	-	0.448	0.577	0.422	0.513
No Glasses	-	-	-	-	0.529	0.385	0.440
No Beards	-	-	-	-	-	0.509	0.440
No Beard No Hats	-	-	-	-	-	-	0.368

Table 29. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the CLIP ViT.

	Female	Male	No Hats	No Glasses	No Beards	No Beard No Hats	No Glasses No Smiles No Ties
Full CelebA	0.513	0.385	0.439	0.294	0.469	0.417	0.255
Female	-	0.203	0.456	0.313	0.312	0.521	0.342
Male	-	-	0.371	0.286	0.238	0.201	0.306
No Hats	-	-	-	0.322	0.411	0.378	0.302
No Glasses	-	-	-	-	0.459	0.427	0.286
No Beards	-	-	-	-	-	0.301	0.324
No Beard No Hats	-	-	-	-	-	-	0.236

Table 30. The full results for the alignment scores ($\mathcal{A}_{\text{score}}$) of the MAE ViT.