

LARGE: LATENT-BASED REGRESSION THROUGH GAN SEMANTICS

Anonymous authors

Paper under double-blind review

ABSTRACT

We propose a novel method for solving regression tasks using few-shot or weak supervision. At the core of our method is the fundamental observation that GANs are incredibly successful at encoding semantic information within their latent space, even in a completely unsupervised setting. For modern generative frameworks, this semantic encoding manifests as smooth, linear directions which affect image attributes in a disentangled manner. These directions have been widely used in GAN-based image editing. We show that such directions are not only linear, but that the magnitude of change induced on the respective attribute is approximately linear with respect to the distance traveled along them. By leveraging this observation, our method turns a pre-trained GAN into a regression model, using as few as two labeled samples. This enables solving regression tasks on datasets and attributes which are difficult to produce quality supervision for. Additionally, we show that the same latent-distances can be used to sort collections of images by the strength of given attributes, even in the absence of explicit supervision. Extensive experimental evaluations demonstrate that our method can be applied across a wide range of domains, leverage multiple latent direction discovery frameworks, and achieve state-of-the-art results in few-shot and low-supervision settings, even when compared to methods designed to tackle a single task.

1 INTRODUCTION

In recent years, Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) have been at the forefront of deep learning research. GANs revolutionized countless generative tasks, such as unconditional image synthesis (Karras et al., 2020b; Brock et al., 2018), cross-domain image-to-image translation (Isola et al., 2017; Zhu et al., 2017a) and super-resolution (Ledig et al., 2017). Beyond generative tasks, numerous works have proposed to use GANs for downstream discriminative objectives, such as classification. Their shared premise is that a generator can synthesize novel samples - often in a controllable manner. These generated samples can then serve as a dataset for training models for downstream tasks. While this approach appears promising on paper (Salimans et al., 2016; Antoniou et al., 2017; Tanaka & Aranha, 2019; Mariani et al., 2018), this simple idea has enjoyed limited success (Cubuk et al., 2020; Ravuri & Vinyals, 2019).

We propose an alternative approach to harnessing the rapid advancement of GANs for downstream tasks. Specifically, we deviate from previous attempts to generate or augment training data. Instead, we focus on extracting information from the incredibly well-behaved latent space of modern GAN architectures, and specifically StyleGAN (Karras et al., 2019; 2020b). The latent spaces of StyleGAN have been studied extensively (Xia et al., 2021), and were shown to be highly semantic, a property which led to their wide use across a range of generative tasks.

In this work, we leverage these semantic properties in order to train regression models. Our key insight is that distances to *semantic hyperplanes* in the latent space are global and discriminative features, accurately gauging the strength of corresponding attributes. The normal vectors of these hyperplanes are the well-known *semantic editing directions* illustrated in Figure 1(a). We show that the latent-space distances can already serve as regression scores for applications where no conventional units are required or exist, *e.g.* ordinal regression tasks. However, we observe that typically these distances are also approximately *linear* with respect to the magnitude of the semantic attribute (see Figure 1(b) for an example). Therefore, they can be calibrated to produce real-world




Figure 1: The two linearities in a GAN’s latent space. (a) A 2D illustration of the **linear** latent directions corresponding to different semantic properties. (b) Yaw angles (Zhou & Gregson, 2020) of generated faces as a function of the distance of their corresponding latent codes from a yaw hyperplane (Shen et al., 2020). There is an approximately **linear** correlation between the two, $R^2 = 0.92$.

predictions, e.g. age in years or head pose in degrees, using linear regression and as few as two labeled samples. This allows us to perform regression in data domains and for semantic attributes where quality supervision is prohibitively difficult to acquire.

Our model, outlined in Figure 2, is thus composed of several steps. First, we learn a disentangled, linear, semantic path for an attribute in the latent space of StyleGAN. Finding such paths is a popular research topic in image editing works (Xia et al., 2021). While editing methods differ in the type of supervision and computation requirements, our approach operates seamlessly with all methods tested. This serves as a promising indication that we can leverage future advancements in editing techniques.

Armed with such a semantic path, we next turn to finding discriminative features which allow us to regress continuous values. We demonstrate that the distance in the latent space from a hyperplane perpendicular to the semantic latent direction is an approximately-linear and smooth representation of the modified property. However, in order to perform regression on real images, the latent code corresponding to a given image is required. For this end, we show that our method is compatible with off-the-shelf *GAN Inversion* encoders, such as pSp (Richardson et al., 2020), e4e (Tov et al., 2021).

Lastly, we observe that the optimal latent space for *GAN inversion* may differ from the optimal space for finding semantic latent directions. In this case, a simple distance from the hyperplane cannot be calculated. To bridge the gap, we learn a task-specific mapping between distances in the two spaces using only latent-space considerations, without any additional supervision. This mapping method is outlined in subsection 3.2. By following these steps, our method distills the strength of any semantic property into a single scalar, using only a pretrained generator and weak supervision.

Through extensive evaluation, we show that our model can produce state-of-the-art results on few-shot learning tasks such as pose and age estimation, without *any* direct supervision on other domains, and that it can even match or outperform fully-supervised methods designed for specific tasks and trained on tens of thousands of samples. Where no supervision is available, we show that our model produces meaningful scores by demonstrating its applicability to the tasks of ordinal regression and sorting collections of images by the strength of a semantic property.

In summary, our contributions are:

- The observation that commonly, latent-space distances are approximately linearly correlated with the magnitude of semantic properties in an image.
- A scheme for converting a pretrained generator and a semantic latent-direction into a state-of-the-art few-shot regressive model.
- A new approach to analysing layer-importance and mapping semantic distances between the latent spaces of a GAN.




Figure 2: Outline of our proposed regression pipeline. An image I_i is inverted into a latent-code w_i . The distance d_i of the code w_i from a semantic hyperplane \vec{P} is calculated. Finally, d_i is input to the simple regression model which outputs the magnitude of the semantic attribute y_i of the image.

2 RELATED WORK

Latent Space of GANs: Recently, understanding and controlling the latent representation of pre-trained GANs has attracted considerable attention. Notably, it has been shown that StyleGAN (Karras et al., 2019; 2020b;a) creates a disentangled, smooth and semantically rich latent space. Many recent works have proposed methods to interpret the semantics encoded in this latent space and apply them to image editing (Shen et al., 2020; Härkönen et al., 2020; Abdal et al., 2020; Wu et al., 2020b; Patashnik et al., 2021; Spingarn-Eliezer et al., 2020; Voynov & Babenko, 2020). Such methods typically identify a **linear** latent direction, which when traversed along, starting from an initial latent code, causes a gradual change in a single semantic property of the corresponding generated image.

To edit real images, one must first obtain the latent code from which the pretrained GAN can most accurately reconstruct the original input image. This task has been commonly referred to as *GAN Inversion* and has been tackled by numerous recent works (Richardson et al., 2020; Tov et al., 2021; Alaluf et al., 2021; Zhu et al., 2020; Abdal et al., 2019; Karras et al., 2020b). For a thorough introduction to these subjects, we refer the reader to a recent survey (Xia et al., 2021).

GANs for discriminative tasks: Several works sought to leverage the advancement of GANs for discriminative tasks. Arguably, the most straight-forward approach is to simply train a generator and use it to synthesize labeled samples. Such samples are inherently labeled when the generator performs either image-to-image translation between domains or class-conditioned synthesis. Following this approach, some works achieved competitive results using completely generated datasets (Tanaka & Aranha, 2019; Marriott et al., 2020). Others enriched a real dataset to improve performance in the low-data domain (Zhu et al., 2017b; Frid-Adar et al., 2018; Antoniou et al., 2017; Wang et al., 2018) or on biased or unbalanced data (Mariani et al., 2018; Hochberg et al., 2021; Ramaswamy et al., 2020; Grover et al., 2019; Sharmanska et al., 2020). When applied to the ImageNet classification task, some works (Ravuri & Vinyals, 2019; Shmelkov et al., 2018) demonstrated that performance sees only modest improvement when enriching the real set, and becomes poor when replacing it.

While these works cover a diverse collection of settings and approaches, they all used GANs in order to eventually generate more data **for training**. In the context of discriminative tasks, several recent methods have proposed to utilize GANs for additional purposes. Lang et al. (2021) used StyleGAN to visualize counterfactual examples for explaining a pretrained classifier’s predictions. Chai et al. (2021) used style-mixing in the fine-layers of StyleGAN to generate augmentations that are ensembled together at test-time. Most related to our work is the representation learning framework GHFeat (Xu et al., 2020). In their work, Xu *et al.* train an encoder for *GAN Inversion* into the latent space of a pretrained StyleGAN and demonstrate that the visual features learned by this encoder can be used to train a variety of models for downstream tasks in a fully supervised manner.

In contrast, we build on recent progress in the study of GANs’ latent space and observe that *distances* within this space can already serve as one-dimensional discriminative features. Our method can leverage these features for simple discriminative tasks, such as sorting, using only weak supervision. With just a few labeled samples, as little as two, we are able to regress real world values. In settings where both our models are applicable, we compare our method with GHFeat and find that we can obtain more accurate results using significantly less data.

Few-shot regression: In the context of visual media, few-shot methods have been widely studied for classification (Snell et al., 2017; Vinyals et al., 2016; Sung et al., 2018), object detection (Wu et al.,

2020a; Fan et al., 2020) and segmentation (Endo & Kanamori, 2021; Rakelly et al., 2018). Relatively few image-related few-shot models have stepped beyond these bounds. They tackle image-to-image translation (Liu et al., 2019), super resolution (Soh et al., 2020), motion prediction (Gui et al., 2018), and re-identification (Wu et al., 2019).

In particular, this scarcity holds for regression tasks, where a majority of research focused on pose estimation using strong task-specific priors (Xiao & Marlet, 2020) or applying meta-learning ideas to keypoint extraction and gaze estimation frameworks (Tseng et al., 2019; Park et al., 2019). For a thorough overview, we refer the reader to a recent survey (Wang et al., 2020). These works, however, all deal with specific tasks and adapt them across domains in a few-shot manner. Their use of task-specific supervision, priors and architectures prevents them from easily being extended to new objectives. In contrast, our method can be easily generalized across domains and different regression goals, including cases where collecting task-specific supervision from other domains would be difficult or outright impossible.

3 METHOD

A shared premise in GAN-based editing works is that single, semantic image properties can be manipulated through modifications of the latent codes used to generate the image. These modifications are conducted by discovering appropriate **global, linear** steering directions within the latent space of the generator (Figure 1(a)). Such directions take the form of vectors, \vec{n} , which can be used to induce a semantic change in *any* code \vec{w} (*i.e.* they are **global**) through **linear** addition: $\vec{w}' = \vec{w} + \alpha\vec{n}$, where α is a scalar that controls the strength of the modification.

Studying these latent directions, we make an important observation – not only are they linear, but oftentimes the magnitude of their effect is linear with respect to α . Traversing one unit of distance across \vec{n} causes a roughly fixed-size change in the property in question (Figure 1(b)). In our work, we rely on this observation and use the distance along a linear latent semantic direction to regress a corresponding property in a given image. Experiments demonstrating the linearity of the distance-features are provided in the appendix.

To regress a single property in a real image I_i , we thus require two components: a semantic latent direction corresponding to the property in question \vec{n} , and the representation of the image in the latent-space of a pretrained GAN: \vec{w}_i . As previously discussed, semantic directions can be discovered using an array of editing methods (Shen et al., 2020; Shen & Zhou, 2020; Patashnik et al., 2021). An alternative approach to considering such directions is to view them as normals to a hyperplane \vec{P} partitioning the latent space: $\vec{n} \cdot \vec{w} + b = 0$. Now, the magnitude of the property in the image is the distance of its latent representation to the hyperplane:

$$d = \text{dist}(\vec{w}_i, \vec{P}) = \vec{w}_i \cdot \vec{n} + b. \quad (1)$$

Some editing methods produce a value for the intercept b , while others do not. When b is unknown, we set it arbitrarily to $b = 0$. As b modifies all distances by a constant factor, dropping it has no effect on any sorting applications, and it can be effectively determined during calibration when training a linear regressor on real data. When b is known, the distance of 0 carries a special semantic meaning. For example, when considering head pose it corresponds to a frontal face.

While calculating the distance is a simple algebraic operation and requires no supervision, this does not imply that the method is unsupervised. Supervision is dictated by the methods used for finding the semantic latent direction and for performing GAN Inversion. These methods typically require weak or indirect supervision. InterFaceGAN (Shen et al., 2020), for example, requires only binary labeling of an attribute. In the head pose example, such annotation would amount to left/right pose information, rather than an explicit yaw angle, which is significantly harder to annotate. The GAN itself is trained in a completely unsupervised manner.

3.1 INVERSION

A requirement of our approach is the ability to invert an image into the latent space of the GAN. Specifically, with StyleGAN, there exist numerous latent spaces that have been considered by previous works. We follow the common approach of inverting an image into the $\mathcal{W}+$ space, first introduced by Abdal et al. (2019). For an inversion method, we follow the encoder-based scheme. Our method works

seamlessly with multiple off-the-shelf encoders, but shows improved performance when utilizing e4e (Tov et al., 2021), which was designed to produce latent codes that are highly editable. We provide a comparison of inversion methods in the appendix.

3.2 BRIDGING THE GAP BETWEEN SPACES

A topic of ongoing research in the field of GAN-based image manipulation is the choice of optimal latent spaces for inversion (Tov et al., 2021) and editing (Wu et al., 2020b). In StyleGAN, the generator supports four latent spaces: \mathcal{Z} , \mathcal{W} , $\mathcal{W}+$ and \mathcal{S} . \mathcal{Z} denotes the usual Gaussian prior, while the others denote increasingly complex spaces obtained by passing through StyleGAN’s mapping network (\mathcal{W}), by assigning different \mathcal{W} codes to different layers of the GAN ($\mathcal{W}+$) or by applying affine transformations to the codes (\mathcal{S}).

A common challenge in StyleGAN editing tasks is that the latent space most often used to identify semantic latent directions, \mathcal{W} , is not expressive enough to support accurate reconstructions of images. Conversely, the $\mathcal{W}+$ space, allows for accurate reconstructions - but behaves poorly under \mathcal{W} -based transformations (Tov et al., 2021). One reason for this behaviour is that in the $\mathcal{W}+$ space, different codes affect different layers which in turn affect different semantic properties of the generated image. For example, pose is controlled by early layers of the network while colors are largely controlled by later layers (Karras et al., 2019; Yang et al., 2021). While applying a \mathcal{W} space editing direction equally to all layer codes may still modify a desired property, not all layer modifications are required, or even affect the property at all. For such layers, the distance from the hyperplane is *irrelevant* to the magnitude of the semantic property. A naïve sum over all layer distances would therefore correlate poorly with the strength of the property. A natural question arising is then - which layers *are* relevant? We answer this question by learning an importance score for each layer in an unsupervised manner.

To do so, we sample a random latent code $\vec{w} \in \mathcal{W}$ in the same space as the semantic hyperplane. We edit the sampled code and obtain $\vec{w}_e = \vec{w} + \alpha \vec{n}$ which is then used to generate an edited image $I_e = G(\vec{w}_e)$. We map the original \vec{w} to a code $\vec{w}^+ \in \mathcal{W}+$ by duplicating it, once for each layer. Finally, we set up a direct optimization scheme where we attempt to modify the mapped code, \vec{w}^+ such that it can be used to generate the edited image, i.e. we aim to solve:

$$\vec{w}^* = \operatorname{argmin}_{\vec{w}^+} \|G(\vec{w}^+) - I_e\|^2. \quad (2)$$

Note that \vec{w}_e is an optimal solution for Eqn. (2). However, it is not the only solution. When optimizing \vec{w}^+ to solve the equation, layers which are *irrelevant* to the property will hardly change from their initialization. The per-layer magnitude of change in the optimized code is therefore an intuitive measure of the importance of each layer towards a semantic attribute. We solve Eqn. (2) for multiple initial codes using gradient descent and track the mean magnitude of the gradients along different layers. In order to avoid spurious results due to different layers having different gradient scales, we normalize gradients by their values when optimizing between unrelated-images. The normalized gradient magnitudes serve as per-layer importance scores, $\{S_i\}_{i=1}^L$, and are used to calculate a weighted sum of hyperplane distances:

$$d_{\mathcal{W}+} = \operatorname{dist}_{\mathcal{W}+}(\vec{w}^+, \vec{P}_{\mathcal{W}}) = \sum_{i=1}^L S_i \cdot \operatorname{dist}_{\mathcal{W}}(\vec{w}_i^+, \vec{P}_{\mathcal{W}}). \quad (3)$$

We use $d_{\mathcal{W}+}$ as an effective distance between $\vec{w}^+ \in \mathcal{W}+$ and a hyperplane $\vec{P}_{\mathcal{W}}$ given in \mathcal{W} .

3.3 CALIBRATION

In many cases, one expects a regression model to output a prediction in “real-world” values, such as head pose in degrees. As presented, the latent-distances are uncalibrated. One can ask, given a sample for which $d = 3$, what is the actual head pose in degrees? At this point, we turn to our prior observation: the magnitude of change in a property is roughly a linear function of the distance. Hence, to calibrate the distances to actual real-world values, we need only train a simple linear regression model with one feature per sample - the distance to the hyperplane. Such a linear model simply takes the form $y = a \cdot d + b$, where y is the calibrated property prediction, d the latent-space distance, and a, b are learned parameters which can be determined with as few as two sampled points.

Once trained, the model gets the distance from the hyperplane as input and predicts the real-world interpretable value. Finally, the entire pipeline produces a few-shot image-property regression model.




Figure 3: Comparisons to hyperplane-distance baselines operating in different feature spaces. Our method outperforms all baselines, indicating GAN-space distances are more semantically meaningful.

4 EXPERIMENTS

We demonstrate our method on several domains and properties. When performing quantitative evaluation, we annotate unlabeled datasets using pretrained networks (specifically WHENet (Zhou & Gregson, 2020) and DEX (Rothe et al., 2015)). Human face image experiments use the official StyleGAN2 pretrained on FFHQ (Karras et al., 2019). Age experiments use the CACD (Chen et al., 2014) and CelebA-HQ (Karras et al., 2017) datasets, while all other evaluations are conducted on CelebA-HQ. Leaf image experiments use the Plant-Village dataset (Hughes et al., 2015). Cat image experiments use the official StyleGAN-ADA (Karras et al., 2020a) model trained on AFHQ (Choi et al., 2020), and are evaluated on the test-split of the same set. Experiments on additional domains are provided in the appendix. In all cases, the GAN was trained in a completely unsupervised manner, without any labels. All results are shown on real images. Where feasible, we report the mean and standard deviation over a thousand repetitions using randomly samples subsets of the data.

4.1 FEATURE SPACE COMPARISONS

We start by demonstrating that distances in the latent space of the GAN are more semantically meaningful and better behaved than equivalent distances in alternative feature spaces. In particular, we compare our method with multiple baselines operating in a similar manner to the InterFaceGAN (Shen et al., 2020) approach. First, a large collection of binary-tagged images are acquired, along with their representation in each chosen feature space. Then, we train an SVM in the given feature space using the binary labels, providing us with a separating hyperplane matching the semantic attribute described by the labels. Finally, the distance to the hyperplane within the given feature space is used as a discriminative feature for training a linear regressor.

The simplest baseline considers the pixel space of the image as a feature space. Other baselines use feature spaces learned by deep neural networks, varying in architectures and tasks. Specifically, we consider the feature spaces learned by: GHFeat (Xu et al., 2020), SwaV (Caron et al., 2020) - an unsupervised representation learning approach, a ResNet101-2W ImageNet classifier (Zagoruyko & Komodakis, 2016) ("ImageNet") and a face recognition network (Deng et al., 2019) ("ID"). Lastly, we compare to the feature space learned by a ResNet18 (He et al., 2016) model trained to provide binary classification of the images for the *same* semantic attribute ("Binary-cls").

As can be seen in Figure 3, for all experiments, our model outperforms the baselines, showing that distances to semantic boundaries within the latent space of a pretrained generator are more semantically meaningful, and serve as better discriminative features for *linear* regression than distances in alternative feature spaces. Furthermore, these results demonstrate that the idea of utilizing a one-dimensional distance metric in some learned feature space is not universal, but relies upon the extensive semantic knowledge encapsulated by the GAN. Indeed, for some configurations, the baselines perform worse than simply returning the mean of the training set’s distribution.

4.2 CALIBRATED RESULTS

Having established that the latent space of StyleGAN is a source of semantically meaningful feature representations, we turn to evaluating the strength of these features with respect to prior works. Towards this end, we demonstrate the performance of our model on a set of regression tasks: pose and age estimation for human faces. Additionally, pose estimation for cars is provided in the appendix.




Figure 4: Comparisons to alternative models as a function of the number of labeled images used in training. In (a) we compare to GHFeat, FSA, and SSV on the CelebA-HQ dataset. In (b) we compare to GHFeat and CORAL on the CACD dataset.

We train and compare models over a wide range of supervision settings, showing that our approach achieves state-of-the-art performance in the few-shot domain as well as surprisingly competitive results on larger datasets, even when compared to methods designed for a single, specific task.

For each attribute we evaluate against a set of methods designed for the specific task, often utilizing a task specific architecture. Furthermore, we compare against GHFeat (Xu et al., 2020), which proposed the use of a different set of GAN-inspired features, and is thus comparable in all scenarios. Comparisons against GHFeat are conducted by training a linear regression model directly on their features. Additionally, we compare to the distance-based GHFeat variant introduced in Subsection 4.1 (“GHFeat SVM”). We find that this variant performs better than the original in few-shot scenarios.

Human Pose: We evaluate the performance of our model on the task of predicting human head poses, and specifically yaw. The semantic direction used by our method was extracted using InterFaceGAN (Shen et al., 2020). In addition to GHFeat, we compare to two dedicated pose estimation methods: FSA (Yang et al., 2019), a fully supervised method and SSV (Mustikovela et al., 2020), a method which learns pose estimations in a self-supervised manner and then calibrates them to dataset-specific values with few labeled samples. In this sense, their method can also be regarded as a few-shot approach. In Figure 4(a) we show the mean absolute errors (MAE) of yaw estimation using the outlined approaches. Our method consistently outperforms all methods when presented with limited supervision, and remains competitive up to a thousand labeled samples. Of particular note is the fact that our model displays better performance than SSV, indicating that the latent space of a pretrained GAN encodes more meaningful pose information than comparable self-supervised methods which were designed and trained specifically to extract pose.

Human Age: We evaluate our performance on the task of human age estimation. The age editing direction used by our method was discovered through natural language descriptions using StyleCLIP (Patashnik et al., 2021). Specifically, the boundary was extracted with the prompts “old face” and “young face”. In addition to GHFeat, we compare to CORAL (Cao et al., 2020). Figure 4(b) shows the MAE on age estimation in years. Our method outperforms the alternatives when presented with limited data, and is only outpaced by CORAL when provided with tens of thousands of samples. Specifically, our 20-sample model obtains a MAE of 7.46, in line with CORAL’s 20K-sample result – 7.59. Note that, unlike pose, age is a property that does not manifest linearly in pixel-space. The visual difference between a person at ages 8 and 13 is much greater than the difference between 30 and 35. However, latent-space linearity still holds in this case. Furthermore, this experiment demonstrates that textually described boundaries are suitable for regression - opening a path to few-shot regression of many properties that can be reasonably described through natural language.

As verified by the experiments, our method achieves state-of-the-art results in the few-shot domain, even when compared to models designed for specific regression tasks. Additionally, our results indicate that the relationship between the latent space distance-features and real-world property values can indeed be well approximated through a *linear* mapping.

4.3 UNCALIBRATED RESULTS

We next discuss domains and attributes for which continuous supervision is not available. Even in such cases, our method can still produce meaningful scores which describe the magnitude of a

semantic property in the image. Without supervision, the scores cannot be calibrated to any real-world human-interpretable value. Nevertheless, uncalibrated scores are still useful for applications such as sorting and ordinal regression. Here, we demonstrate their applicability to the task of sorting a collection of images by a given property – *e.g.* according to how happy the person in the image is. An ordinal regression experiment is provided in the appendix.

In Figure 5 we show results for sorting a set of face images according to properties described through textual prompts, using distances to boundaries discovered through StyleCLIP’s global mapping approach. We sort the same randomly sampled collection according to four distinct attributes: expression, hair color, hair length, and amount of makeup.

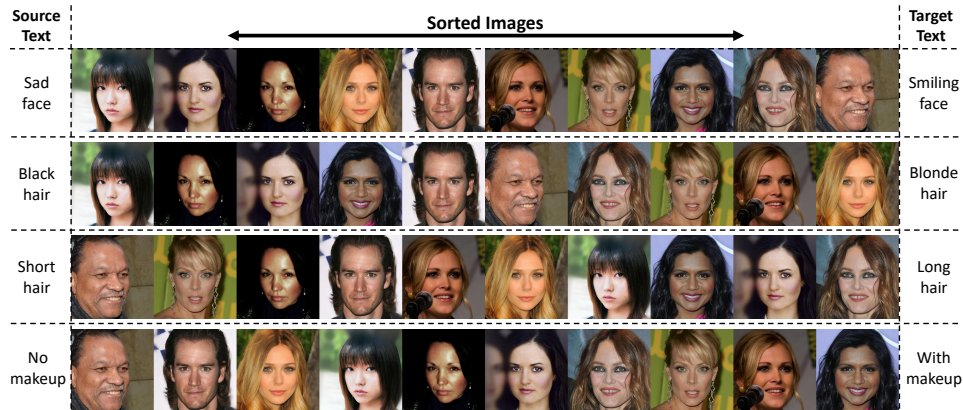


Figure 5: Sorting images according to textual descriptions of semantic properties. In each row, we sort the same set of randomly sampled CelebA-HQ images according to their distance from a text-based editing boundary extracted by StyleCLIP. Each row’s editing direction is induced by the source text (left) and the target text (right).




Figure 6: Sorting images from the Plant-Village dataset using a semantic direction extracted by InterFaceGAN. In each row, we sort randomly sampled sets containing images labeled as either healthy or sick. We separate rows by type of disease to facilitate visual comparisons.

In Figure 6 we show the results of sorting a collection of leaf images corresponding to a ‘sick or healthy’ direction, using a boundary extracted with InterFaceGAN (Shen et al., 2020) based on binary annotations. Our method successfully turns these binary annotations into continuous values which allow us to determine which leaves are sicker than others. For example, in the last row, the number of “black spots” decreases gradually while moving from the most-sick leaf towards the healthy leaves.

In Figure 7 we show results on sorting collections of cats randomly sampled from the AFHQ dataset (Choi et al., 2020), using semantic directions discovered in an unsupervised manner with SeFA (Shen & Zhou, 2020). Our method extends seamlessly to these additional domains and semantic direction discovery approaches.

As can be seen, the ordered results largely align with human expectation. In order to verify this claim we conduct a user study on the human and cat face domains. The full details of the baselines and the

experimental setup are provided in the appendix. Results are summarized in Table 1. As verified by the study, our model learns to regress more consistent scores for images across both domains.

As demonstrated through our experiments, our approach learns to regress meaningful, uncalibrated values across multiple domains and using a wide range of latent-direction discovery methods.



Figure 7: Sorting images from AFHQ-cat using semantic directions extracted by SeFA.

Table 1: User study results for sorting quality. For each attribute we report the percent of responders which preferred the sorting induced by each of three sorting methods. Our method consistently provides an order which is more consistent with human preference. See the appendix for more details.

Attribute	Ours	CLIP	Random	Attribute	Ours	SSV	Random
Hair color	73.55%	25.81%	0.65%	Yaw	66.67%	28.99%	4.35%
Makeup	70.53%	18.25%	11.23%	Pitch	82.50%	7.50%	10.00%
Expression	53.45%	43.27%	3.27%	Average	75.71%	16.71%	7.58%
Hair length	84.73%	2.18%	13.09%				
Average	70.56%	22.38%	7.06%				

(a) Human faces

(b) Cats

5 DISCUSSION

We have presented a method for leveraging the semantic structure of a pre-trained GAN for weakly-supervised and few-shot regression tasks. Our method builds upon extensive prior works which explored the latent space of StyleGAN. Some of these works found semantic linear latent directions and applied them to image editing. Others directly utilized latent codes as features for regression tasks. However, we are the first to note that insights from both approaches can be combined. We have shown that distances to semantic hyperplanes can serve as global, descriptive features. Additionally, we demonstrated that they are typically linearly correlated with the magnitude of change in semantic attributes. We then used these features to train simple linear regression models, with as few as two labeled samples, and obtained state-of-the-art results in the few-shot domain.

Our work focused on the state of the art in general image synthesis - StyleGAN. However, it is not the only generative network to exhibit an editable latent space. Alternative networks such as BigGAN (Brock et al., 2018) similarly support editing operations (Jahani et al., 2019; Spingarn-Eliezer et al., 2020), and may allow for regression on attributes or datasets which StyleGAN struggles with.

While we have observed linearity in all semantic attributes which we tested, this property is unlikely to hold universally. Indeed, for some attributes, even finding a disentangled latent direction is infeasible. Furthermore, similar to other methods which rely on StyleGAN (Lang et al., 2021; Chai et al., 2021), our method obtains better results when operating on images within the domain used to train the GAN. This limitation stems in part from the inability of current *GAN Inversion* methods to reconstruct out-of-domain images while preserving latent semantics. However, as this semantic-preserving out-of-domain inversion task is of interest to the community at large, we are looking forward to see this barrier overcome.

We hope that our work can inspire others to consider the latent space of GANs as a source of semantically-rich supervision which can be leveraged to tackle a wide range of downstream tasks.

6 ETHICS STATEMENT

Our model consists of a general method for performing regression on images. As such, its impact is dependent on the tasks for which it is used. Such tasks could have a wide range of positive benefits in numerous computer vision related fields. For example, the ability to quantify levels of disease in plants, and perhaps in other domains, may be of benefit in the agricultural and healthcare fields. Enabling few-shot regression may further assist with 'democratizing' neural networks, in the sense that the method could enable smaller businesses or research groups in performing regression in scenarios where labeled data may be out of their means.

On the other hand, our model could be used in applications which violate privacy, for example by facilitating the collection of information such as age or sentiments from individual photographs.

Furthermore, our model is sensitive to the same biases found in the data used to train the GAN - and in the case of natural-language based regression, also to the biases present in CLIP. As such, it may assist in perpetuating biases such as gender norms (as seen in the makeup sorting experiment) or racial discrimination (*e.g.* Asian descent is correlated with age in the data set, and this is reflected in model predictions).

7 REPRODUCIBILITY

We provide the URL to an anonymized Github repository containing our source code: <https://github.com/iclr2022-paper1352/iclr2022-paper-1352>. Additionally, we have made an effort throughout the paper towards making the results reliable and reproducible. Specifically, as noted in Section 4, whenever feasible we repeated experiments a thousand times and report mean and standard deviation in Figures 3, 4, 8 to 10, 13 and 26. Additionally, we report detailed information on subtle aspects of the model and data in Section D.

REFERENCES

- Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE international conference on computer vision*, pp. 4432–4441, 2019. 3, 4
- Rameen Abdal, Peihao Zhu, Niloy Mitra, and Peter Wonka. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows, 2020. 3
- Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Restyle: A residual-based stylegan encoder via iterative refinement. *arXiv preprint arXiv:2104.02699*, 2021. 3, 15, 16, 33
- Antreas Antoniou, Amos Storkey, and Harrison Edwards. Data augmentation generative adversarial networks. *arXiv preprint arXiv:1711.04340*, 2017. 1, 3
- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 1, 9
- Wenzhi Cao, Vahid Mirjalili, and Sebastian Raschka. Rank consistent ordinal regression for neural networks with application to age estimation. *Pattern Recognition Letters*, 140:325–331, 2020. 7
- Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*, 2020. 6, 31
- Lucy Chai, Jun-Yan Zhu, Eli Shechtman, Phillip Isola, and Richard Zhang. Ensembling with deep generative views. In *CVPR*, 2021. 3, 9
- Bor-Chun Chen, Chu-Song Chen, and Winston H. Hsu. Cross-age reference coding for age-invariant face recognition and retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. 6, 32, 33

- Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 6, 8, 20, 21, 22, 24, 33
- Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020. 1
- Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4690–4699, 2019. 6, 31
- Yuki Endo and Yoshihiro Kanamori. Few-shot semantic image synthesis using stylegan prior. *arXiv preprint arXiv:2103.14877*, 2021. 4
- Qi Fan, W. Zhuo, and Yu-Wing Tai. Few-shot object detection with attention-rpn and multi-relation detector. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4012–4021, 2020. 4
- Maayan Frid-Adar, Idit Diamant, Eyal Klang, Michal Amitai, Jacob Goldberger, and Hayit Greenspan. Gan-based synthetic medical image augmentation for increased cnn performance in liver lesion classification. *Neurocomputing*, 321:321–331, 2018. 3
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014. 1
- Aditya Grover, Kristy Choi, Rui Shu, and Stefano Ermon. Fair generative modeling via weak supervision. 2019. 3
- Liang-Yan Gui, Yu-Xiong Wang, Deva Ramanan, and José MF Moura. Few-shot human motion prediction via meta-learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 432–450, 2018. 4
- Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *arXiv preprint arXiv:2004.02546*, 2020. 3
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016. 6
- Dana Cohen Hochberg, Raja Giryes, and Hayit Greenspan. Style encoding for class-specific image generation. In *Medical Imaging 2021: Image Processing*, volume 11596, pp. 1159631. International Society for Optics and Photonics, 2021. 3
- David Hughes, Marcel Salathé, et al. An open access repository of images on plant health to enable the development of mobile disease diagnostics. *arXiv preprint arXiv:1511.08060*, 2015. 6, 20, 22, 23, 33
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134, 2017. 1
- Ali Jahanian, Lucy Chai, and Phillip Isola. On the “steerability” of generative adversarial networks. *arXiv preprint arXiv:1907.07171*, 2019. 9
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 6, 15, 16, 19, 20, 32, 33
- Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4401–4410, 2019. 1, 3, 5, 6, 32, 33

- Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Proc. NeurIPS*, 2020a. [3](#), [6](#), [20](#), [33](#)
- 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*, pp. 8110–8119, 2020b. [1](#), [3](#), [19](#), [33](#)
- Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pp. 554–561, 2013. [19](#), [33](#)
- Oran Lang, Yossi Gandelsman, Michal Yarom, Yoav Wald, Gal Elidan, Avinatan Hassidim, William T Freeman, Phillip Isola, Amir Globerson, Michal Irani, et al. Explaining in style: Training a gan to explain a classifier in stylespace. *arXiv preprint arXiv:2104.13369*, 2021. [3](#), [9](#)
- Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4681–4690, 2017. [1](#)
- Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10551–10560, 2019. [4](#)
- Giovanni Mariani, Florian Scheidegger, Roxana Istrate, Costas Bekas, and Cristiano Malossi. Bagan: Data augmentation with balancing gan. *arXiv preprint arXiv:1803.09655*, 2018. [1](#), [3](#)
- Richard T Marriott, Safa Madiouni, Sami Romdhani, Stéphane Gentic, and Liming Chen. An assessment of gans for identity-related applications. In *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–10. IEEE, 2020. [3](#)
- Siva Karthik Mustikovela, Varun Jampani, Shalini De Mello, Sifei Liu, Umar Iqbal, Carsten Rother, and Jan Kautz. Self-supervised viewpoint learning from image collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3971–3981, 2020. [7](#), [24](#), [33](#)
- Seonwook Park, Shalini De Mello, Pavlo Molchanov, Umar Iqbal, Otmar Hilliges, and Jan Kautz. Few-shot adaptive gaze estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9368–9377, 2019. [4](#)
- Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. *arXiv preprint arXiv:2103.17249*, 2021. [3](#), [4](#), [7](#), [20](#), [33](#)
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011. [32](#), [33](#)
- 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. *arXiv preprint arXiv:2103.00020*, 2021. [24](#), [33](#)
- Kate Rakelly, Evan Shelhamer, Trevor Darrell, Alexei A Efros, and Sergey Levine. Few-shot segmentation propagation with guided networks. *arXiv preprint arXiv:1806.07373*, 2018. [4](#)
- Vikram V Ramaswamy, Sunnis SY Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. *arXiv preprint arXiv:2012.01469*, 2020. [3](#)
- Suman Ravuri and Oriol Vinyals. Classification accuracy score for conditional generative models. *arXiv preprint arXiv:1905.10887*, 2019. [1](#), [3](#)

- Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. *arXiv preprint arXiv:2008.00951*, 2020. 2, 3, 15, 16, 33
- Rasmus Rothe, Radu Timofte, and Luc Van Gool. Dex: Deep expectation of apparent age from a single image. In *IEEE International Conference on Computer Vision Workshops (ICCVW)*, December 2015. 6, 33
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *arXiv preprint arXiv:1606.03498*, 2016. 1
- Kim Seonghyeon. Stylegan2-pytorch. <https://github.com/rosinality/stylegan2-pytorch>, 2020. 33
- Viktoria Sharmanska, Lisa Anne Hendricks, Trevor Darrell, and Novi Quadrianto. Contrastive examples for addressing the tyranny of the majority. *arXiv preprint arXiv:2004.06524*, 2020. 3
- Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. *arXiv preprint arXiv:2007.06600*, 2020. 4, 8, 20, 21, 22, 33
- Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9243–9252, 2020. 2, 3, 4, 6, 7, 8, 19, 20, 22, 23, 33
- Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. How good is my gan? In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 213–229, 2018. 3
- Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175*, 2017. 3
- Jae Woong Soh, Sunwoo Cho, and Nam Ik Cho. Meta-transfer learning for zero-shot super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3516–3525, 2020. 4
- Nurit Spingarn-Eliezer, Ron Banner, and Tomer Michaeli. Gan steerability without optimization. *arXiv preprint arXiv:2012.05328*, 2020. 3, 9
- Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1199–1208, 2018. 3
- Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pp. 6105–6114. PMLR, 2019. 19
- Fabio Henrique Kiyoyiti dos Santos Tanaka and Claus Aranha. Data augmentation using gans. *arXiv preprint arXiv:1904.09135*, 2019. 1, 3
- Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *arXiv preprint arXiv:2102.02766*, 2021. 2, 3, 5, 15, 16, 19, 33
- Hung-Yu Tseng, Shalini De Mello, Jonathan Tremblay, Sifei Liu, Stan Birchfield, Ming-Hsuan Yang, and Jan Kautz. Few-shot viewpoint estimation. *arXiv preprint arXiv:1905.04957*, 2019. 4
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. *arXiv preprint arXiv:1606.04080*, 2016. 3
- Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. *arXiv preprint arXiv:2002.03754*, 2020. 3
- Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM Computing Surveys (CSUR)*, 53(3):1–34, 2020. 4
- Yu-Xiong Wang, Ross Girshick, Martial Hebert, and Bharath Hariharan. Low-shot learning from imaginary data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7278–7286, 2018. 3

- Jiaxi Wu, Songtao Liu, Di Huang, and Yunhong Wang. Multi-scale positive sample refinement for few-shot object detection. In *European Conference on Computer Vision*, pp. 456–472. Springer, 2020a. [3](#)
- Lin Wu, Yang Wang, Hongzhi Yin, Meng Wang, and Ling Shao. Few-shot deep adversarial learning for video-based person re-identification. *IEEE Transactions on Image Processing*, 29:1233–1245, 2019. [4](#)
- Zongze Wu, Dani Lischinski, and Eli Shechtman. Stylespace analysis: Disentangled controls for stylegan image generation, 2020b. [3](#), [5](#)
- Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. Gan inversion: A survey. *arXiv preprint arXiv:2101.05278*, 2021. [1](#), [2](#), [3](#)
- Yang Xiao and Renaud Marlet. Few-shot object detection and viewpoint estimation for objects in the wild. In *European Conference on Computer Vision*, pp. 192–210. Springer, 2020. [4](#)
- Yang Xiao, Yuming Du, and Renaud Marlet. Posecontrast: Class-agnostic object viewpoint estimation in the wild with pose-aware contrastive learning. In *ArXiv*, 2021. [19](#), [33](#)
- Yinghao Xu, Yujun Shen, Jiapeng Zhu, Ceyuan Yang, and Bolei Zhou. Generative hierarchical features from synthesizing images. *arXiv e-prints*, pp. arXiv–2007, 2020. [3](#), [6](#), [7](#), [16](#), [31](#), [33](#)
- Ceyuan Yang, Yujun Shen, and Bolei Zhou. Semantic hierarchy emerges in deep generative representations for scene synthesis. *International Journal of Computer Vision*, 129(5):1451–1466, 2021. [5](#)
- Linjie Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. A large-scale car dataset for fine-grained categorization and verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3973–3981, 2015. [19](#), [33](#)
- Tsun-Yi Yang, Yi-Ting Chen, Yen-Yu Lin, and Yung-Yu Chuang. Fsa-net: Learning fine-grained structure aggregation for head pose estimation from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1087–1096, 2019. [7](#), [33](#)
- Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. [19](#)
- Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. [6](#), [31](#)
- Yijun Zhou and James Gregson. Whenet: Real-time fine-grained estimation for wide range head pose. *arXiv preprint arXiv:2005.10353*, 2020. [2](#), [6](#), [33](#)
- Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. *arXiv preprint arXiv:2004.00049*, 2020. [3](#)
- 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*, pp. 2223–2232, 2017a. [1](#)
- Xinyue Zhu, Yifan Liu, Zengchang Qin, and Jiahong Li. Data augmentation in emotion classification using generative adversarial networks. *arXiv preprint arXiv:1711.00648*, 2017b. [3](#)

Appendices

A ABLATION STUDY

We conduct an ablation study on several components of our method. Namely, the use of a linear regression model, our choice of inversion mechanism and the importance of our layer-weighting approach (as detailed in Subsection 3.2 of the core paper).

A.1 EVALUATING LINEARITY

We provide an additional experiment demonstrating the approximately linear correlation of the attribute’s magnitude and the distance-features. Using the same set of labeled images and their distances from a semantic hyperplane, we compare the accuracy of a linear model trained on these data to the accuracy of polynomial models of higher degrees - specifically 2, 3 and 5. Results are shown in Figure 8, As can be seen, for all attributes tested, the linear model is superior when a few labeled samples are provided. Additionally, despite having greater expressive power, the polynomial models do not outpace the linear model even when a thousand labeled samples are provided.

The ability of a linear model to more accurately capture the correlation between the distance-feature and the magnitude of the attribute, implies that it is indeed approximately linear.

A.2 INVERSION COMPARISONS

We evaluate the performance of our model when utilizing different inversion methods in order to obtain the latent code for both train and test images. We compare our results on the human face pose and age estimation tasks using the CelebA-HQ (Karras et al., 2017) dataset. Specifically, we compare four *GAN Inversion* encoder models - pSp (Richardson et al., 2020), e4e (Tov et al., 2021), ReStyle-*psp* and ReStyle-*e4e* (Alaluf et al., 2021) which uses the former encoders in an iterative refinement scheme. In all cases we use the official pre-trained models provided by the authors.

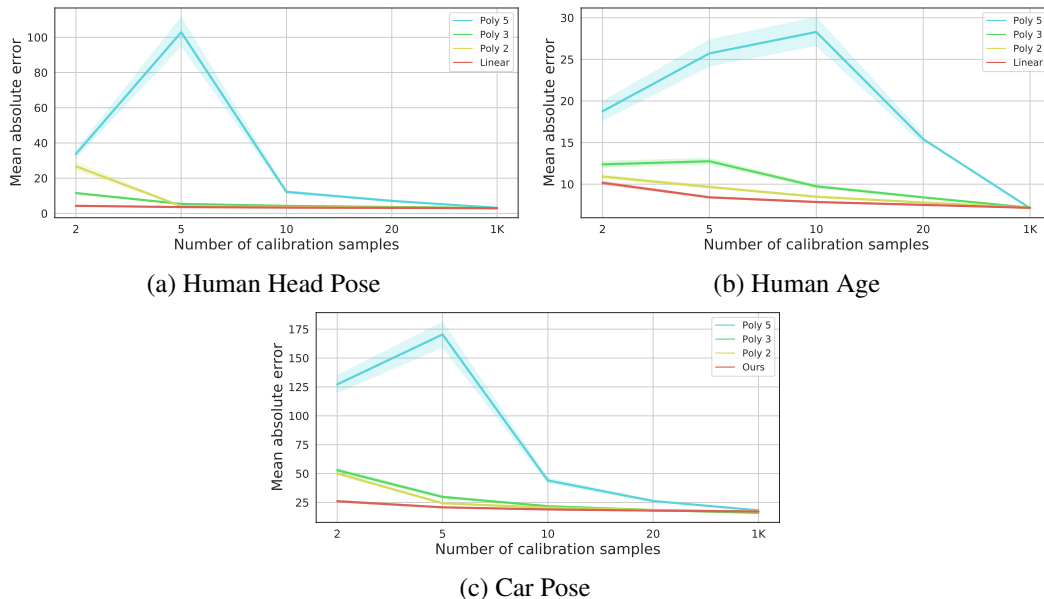


Figure 8: Quantitative comparison of four different choices for the degree of fitted regression functions. In addition to the linear model outlined in our method, we evaluate polynomial degrees of 2, 3 and 5. The linear model outperforms the alternatives when only a few samples are available, and is equivalent for a thousand samples.




Figure 9: Quantitative comparisons of four different *GAN Inversion* encoders: e4e (Tov et al., 2021), pSp (Richardson et al., 2020) and ReStyle e4e and ReStyle (Alaluf et al., 2021) on the CelebA-HQ dataset (Karras et al., 2017).

For reference, we also include the most competitive baseline, GHFeat-SVM which was introduced in Subsection 4.1. As a reminder, this baseline was devised based on the same latent-distance principles but applied in the feature space of GHFeat (Xu et al., 2020). Results are displayed in Figure 9. As can be seen, our method outperforms GHFeat-SVM regardless of the choice of inversion method. Furthermore, e4e (Tov et al., 2021) is consistently superior to other methods, and e4e-based methods are superior to pSp-based methods. Tov et al. (2021) demonstrated that there exists a trade-off between distortion and editability. This tradeoff stems from the ability to invert images into more semantic regions of the latent space. While the e4e encoder tends towards preserving such semantics, pSp is trained with the sole purpose of image reconstruction. As our method relies heavily on such latent space semantics, it is unsurprising to see a consistent, even though minor, advantage towards e4e-based methods.

While superior results are obtained with all inversion methods, we conclude that our method works best with semantic preserving encoders and recommend using such.

A.3 BRIDGING THE GAP

We evaluate the contribution of our per-layer latent direction weighting approach, described in Subsection 3.2 of the main paper.

We compare our proposed approach with three alternatives for computing distances between latent codes in $\mathcal{W}+$ and boundaries found in \mathcal{W} . First, we use a simple model dubbed “All Layers” in which we compute the distance of the latent code in each separate layer to the same \mathcal{W} -space boundary. We then use all such distances (18 in total) as features for the regression model. Second, we consider a model dubbed “Euclidean”, in which we duplicate the boundary over all layers and compute a simple Euclidean distance. Last, we consider a model which uses only the distance along the single layer which provides optimal performance. In the case of poses, for example, this is layer 2. Note that determining the optimal layer in this manner requires a large continuously tagged dataset to evaluate against, which may not be feasible in practical applications. Our own model uses a weighted distance metric where the contribution of each layer is scaled according to our semantic-mapping importance scores determined in an unsupervised manner.

The performance of our method and all alternative is reported in Figure 10. Our weighted model consistently outperforms the other alternatives in low-supervision settings, and achieves similar results to the “All Layers” model with extensive supervision. This shows that our weighted distances accurately reflect the importance of the distance across each layer, without having to rely on the additional supervision required to determine such a weighting with multiple-feature regression.

Beside obtaining superior performance, we next demonstrate that our approach in fact identifies layers that are highly correlated with a semantic attribute. For this end, we first fit a set of linear models for the pose estimation task, using the individual distances along each layer of $\mathcal{W}+$, one at a time. The results and their R^2 coefficients are shown in Figure 12. Note that obtaining accurate correlation scores in this manner requires a large labeled dataset and may hence be unfeasible. Next, in Figure 11 we show the layer importance scores extracted by our unsupervised method. As can be seen, our unsupervised layer scoring approach successfully identifies layers with high correlation to the semantic property.




Figure 10: Comparing several approaches for calculating latent-space distances between codes in $\mathcal{W}+$ and boundaries in \mathcal{W} . The "Euclidean" model calculates the Euclidean distance between the latents and a boundary obtained by replicating the \mathcal{W} boundary along all layers of $\mathcal{W}+$. "All Layers" calculates a per-layer distance and uses all 18 distances as features. "Layer 2" uses only the distance calculated on layer 2 of the latent code, which was experimentally observed to provide the best single-layer results for pose. Finally, our model uses a weighted distance as outlined in Subsection 3.2 of the main paper. As can be seen, our proposed method is superior to other method in the few-shot domain and is matched by "All Layers" only when provided with a thousand labeled samples.




Figure 11: The results of our per-layer importance scores approach as outlined in Subsection 3.2, for the head pose attribute. (a) Un-normalized importance scores, before accounting for the scale of gradients in each layer. (b) Normalized importance scores, after accounting for gradient scales.

A.4 COMPARING AGAINST GENERATING A CONTINUOUSLY-LABELED DATA

As discussed in the paper, a popular approach to using Deep Generative Models for discriminative task is to generate labeled datasets and use them for training. Although our method provides means to perform regression directly in the latent space, it may also be used to generate labeled datasets. After calibrating the model, we can apply it to any latent code, including codes simply sampled from the Gaussian prior. We can thus generate a new dataset from randomly sampled codes, and dictate their labels by the latent regression model.

We perform an experiment where we generate such a dataset for human head pose and train a regression CNN directly on the generated images. To make sure the attribute varies enough in the generated dataset we perform the following process. We sample a latent code $w \in \mathcal{W}$ and a scalar $\alpha \in [-9, 9]$ which was observed to be the maximal range for which the generated image does not degrade in quality. We then edit the sampled latent code by applying $w' = w + \alpha \vec{n}$. Now, we generate the image $I' = G(w')$ and infer the head pose by inputting α to the regression model. For the regression model, we use a model trained with 1000 labeled samples. Repeating this process 45K times, we now possess a continuously-labeled, roughly balanced generated dataset. We train multiple

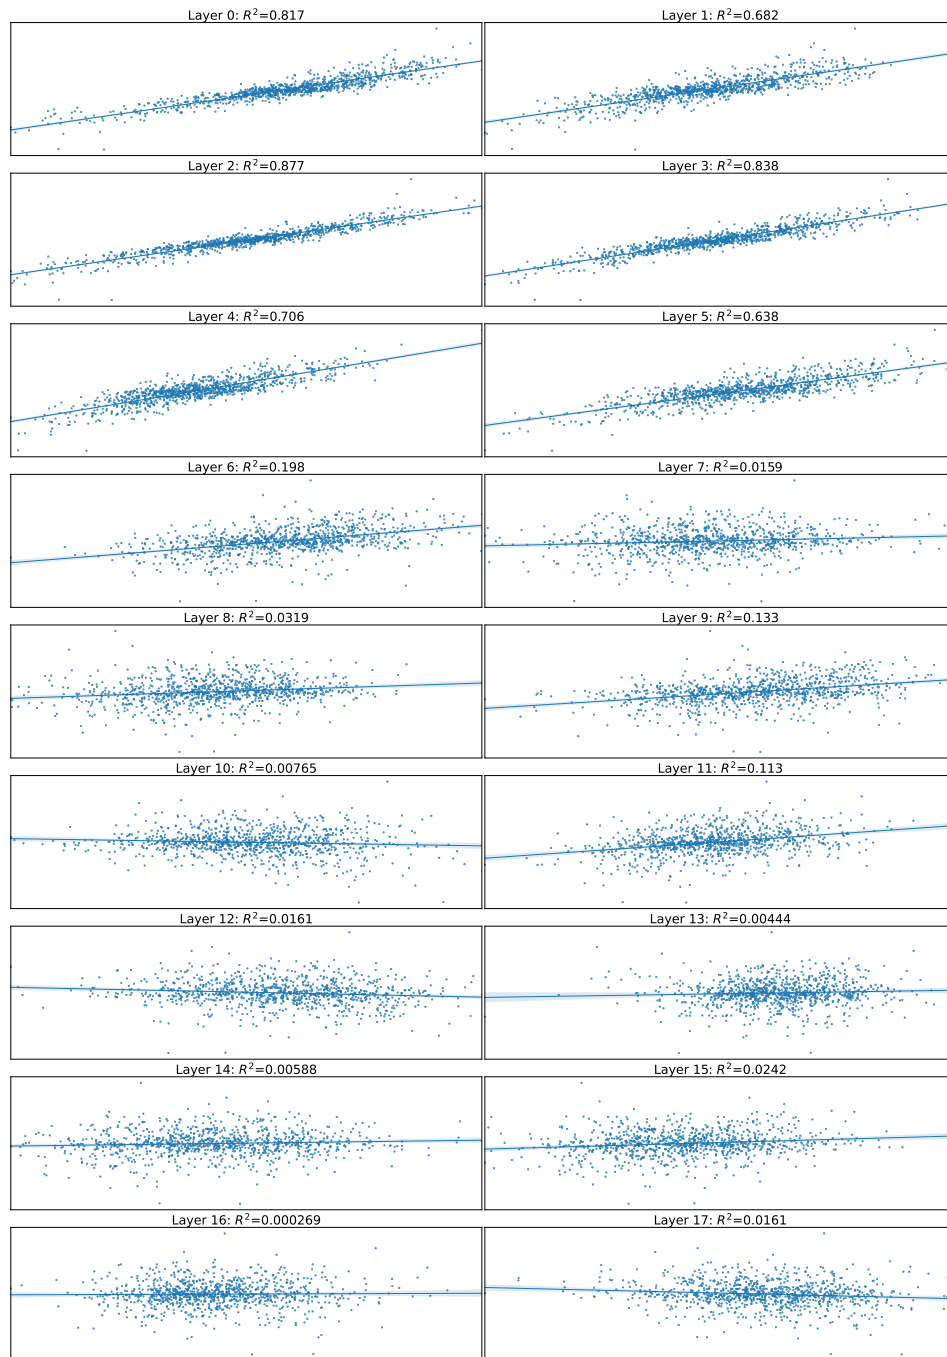


Figure 12: Measuring the linear correlation of yaw angle with distance from hyperplane for each layer separately. In each subplot, the x-axis is the distance of this layer in the latent code from the boundary while the y-axis is the ground truth yaw angle. As can be seen, distance in first layers are better linearly correlated to head pose than last layers.

CNN backbones for the task of regression and test them over the annotated CelebA-HQ test set (Karras et al., 2017). The lowest Mean Average Error, 3.23 ± 3.67 , was obtained with EfficientNet-b3 (Tan & Le, 2019). For comparison, applying our approach with exactly the same set and supervision obtained 2.97 ± 2.76 .

We conclude that it is preferable to apply our method directly to regress test images, rather than generating a labeled set for downstream training. We speculate that a possible explanation for the degradation in performance is the domain gap. In the generated-dataset approach, the classifier is trained on a generated dataset and tested on a real one, without adaptation. On the other hand, in our approach, the *GAN Inversion* encoder may mitigate some of that gap. Additionally, the generator and inversion encoder are trained once per-domain while the latent and CNN regressor as well as the data generation happens once per attribute. As a result, our approach, requiring just the training of a latent simple regression model requires roughly x1000 less time to train.

B ADDITIONAL RESULTS

B.1 CALIBRATED RESULTS - CARS

We repeat the pose experiment of section 4.2 on the car image domain. We compare our model to SSV. Our model utilizes the official StyleGAN2 (Karras et al., 2020b) LSUN Car (Yu et al., 2015) pretrained checkpoint, and the official e4e (Tov et al., 2021) inversion encoder trained on the train split of Stanford Cars (Krause et al., 2013). The semantic boundary was extracted using InterFaceGAN (Shen et al., 2020). SSV was trained as in the original paper, using the CompCars dataset (Yang et al., 2015). Both models were evaluated on the test split of Stanford Cars, with pose labels acquired through Pose Contrast (Xiao et al., 2021). After labeling the test-set images, we discarded all images with yaw angles exceeding 90° in either direction, *i.e.* we evaluated only on images for which $\theta_{yaw} \in [-90^\circ, 90^\circ]$.

The results are shown in Figure 13. Our model outperforms SSV over all tested supervision ranges, indicating that it can generalize well to the car domain.




Figure 13: Pose estimation error comparisons on the Stanford Car (Krause et al., 2013) dataset tagged by Pose Contrast (Xiao et al., 2021), as a function of the number of labeled images used for calibration.

B.2 UNCALIBRATED RESULTS

As discussed in the core paper, our method can be applied to downstream tasks even in the absence of direct supervision. We demonstrated the applicability of our method to an image sorting task. Here, we demonstrate an application to ordinal regression. Specifically, we perform sentiment analysis on facial images, dividing them into four bins that represent different levels of contentment.

To perform ordinal regression, we use our method to calculate an uncalibrated distance score for each image, as described in section 3. Then, we simply divide the range of distances into four bins. As can be seen in Figure 14, even this simple method obtains good results.

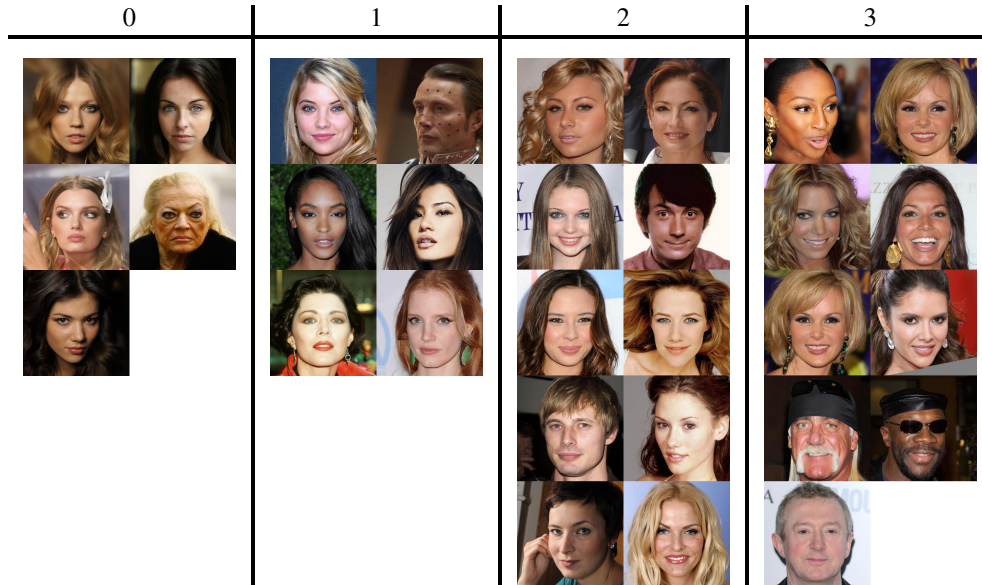


Figure 14: Ordinal regression applied to sentiment analysis using our method. Images are divided into bins, from discontent - 0 to most content - 4. All images were randomly sampled from CelebA-HQ (Karras et al., 2017). Sentiment is measured by distance from a ‘smiling’ semantic boundary, identified by StyleCLIP (Patashnik et al., 2021).

In Figures 15 to 20, we show additional sorting results on cats and dogs, using a StyleGAN-ADA (Karras et al., 2020a) model trained on the AFHQ dataset (Choi et al., 2020). Latent semantics were identified using SeFA (Shen & Zhou, 2020). In Figures 21 to 24, we show additional sorting results on leaves, using a StyleGAN-ADA model trained on the Plant-Village (Hughes et al., 2015) dataset. Latent semantics were identified using InterFaceGAN (Shen et al., 2020).



Figure 15: Sorting images from AFHQ-dog (Choi et al., 2020) using a “fur fluffiness” semantic directions extracted by SeFA (Shen & Zhou, 2020).



Figure 16: Sorting images from AFHQ-dog (Choi et al., 2020) using a “head pitch” semantic directions extracted by SeFA (Shen & Zhou, 2020).



Figure 17: Sorting images from AFHQ-dog (Choi et al., 2020) using a “head yaw” semantic directions extracted by SeFA (Shen & Zhou, 2020).



Figure 18: Sorting images from AFHQ-cat (Choi et al., 2020) using an “age” semantic directions extracted by SeFA (Shen & Zhou, 2020).



Figure 19: Sorting images from AFHQ-cat (Choi et al., 2020) using a “head pitch” semantic directions extracted by SeFA (Shen & Zhou, 2020).



Figure 20: Sorting images from AFHQ-cat (Choi et al., 2020) using a “head yaw” semantic directions extracted by SeFA (Shen & Zhou, 2020).



Figure 21: Sorting images from Plant-Village (Hughes et al., 2015) using sick-to-healthy semantic directions extracted by InterFaceGAN (Shen et al., 2020). To facilitate easy visual comparisons, all sick leaves have the same disease - “Early Blight”.




Figure 22: Sorting images from Plant-Village (Hughes et al., 2015) using sick-to-healthy semantic directions extracted by InterFaceGAN (Shen et al., 2020). To facilitate easy visual comparisons, all sick leaves have the same disease - “Black Rot”.



Figure 23: Sorting images from Plant-Village (Hughes et al., 2015) using sick-to-healthy semantic directions extracted by InterFaceGAN (Shen et al., 2020). To facilitate easy visual comparisons, all sick leaves have the same disease - “Rust”.



Figure 24: Sorting images from Plant-Village (Hughes et al., 2015) using sick-to-healthy semantic directions extracted by InterFaceGAN (Shen et al., 2020). To facilitate easy visual comparisons, all sick leaves have the same disease - “Late Blight”.

C USER STUDY

In this section we provide all the details of our user study.

The user study was conducted through <https://freeonlinesurveys.com/>. The study was performed over a period of 5 days, with a total of 62 different responders. Individual questions saw anywhere from 20 to 62 responses (see Table 2 for exact numbers). All responders were unpaid volunteers which responded (anonymously) to a link shared among colleagues and acquaintances of the authors.

We used a three-alternative forced choice setting. Users were provided with randomly sampled sets of 10 images, sorted in three manners - once using our method, once by randomly assigning an order and once by using a dedicated baseline. For each question, the visual order of the three sorting options was randomized. Users were asked to choose the order that better matches a textual description.

For the human face domain, we used the same textual prompts as Figure 5 in the main paper, with 5 randomly sampled image sets for each prompt. For a baseline, we sorted the images according to their cosine-distances from the same textual prompts directly in the CLIP (Radford et al., 2021) embedding space. For the cat domain we used the pitch and yaw directions displayed in Figure 7 in the main paper and for a baseline we used pose predictions from SSV (Mustikovela et al., 2020) trained on AFHQ-Cat (Choi et al., 2020).

The questions and their associated image sets are shown in Figure 25. In Table 2 we provide the list of the methods used to generate each row of each question (*i.e.* the answer key), along with the number of responders who chose each answer.

Table 2: User study answer key and the number of responders that picked each answer

Question	Ours	CLIP	Random
----- Hair color -----			
1	Bottom (46)	Top (16)	Middle (0)
2	Middle (50)	Top (11)	Bottom (1)
3	Bottom (34)	Top (27)	Middle (1)
4	Middle (53)	Top (9)	Bottom (0)
5	Bottom (45)	Middle (17)	Top (0)
----- Makeup -----			
6	Bottom (45)	Top (8)	Middle (4)
7	Bottom (35)	Top (9)	Middle (13)
8	Middle (36)	Bottom (17)	Top (4)
9	Bottom (41)	Middle (11)	Top (5)
10	Top (44)	Middle (7)	Bottom (6)
----- Expression -----			
11	Bottom (40)	Top (13)	Middle (2)
12	Bottom (46)	Top (9)	Middle (0)
13	Bottom (38)	Top (15)	Middle (2)
14	Top (19)	Bottom (35)	Middle (1)
15	Bottom (4)	Middle (47)	Top (4)
----- Hair length -----			
16	Top (52)	Bottom (1)	Middle (2)
17	Middle (46)	Top (1)	Bottom (8)
18	Top (33)	Middle (0)	Bottom (22)
19	Middle (50)	Bottom (1)	Top (4)
20	Middle (52)	Top (3)	Bottom (0)
----- Pitch -----			
21	Bottom (4)	Middle (18)	Top (1)
22	Top (22)	Middle (0)	Bottom (1)
23	Top (20)	Middle (2)	Bottom (1)
----- Yaw -----			
24	Bottom (17)	Middle (0)	Top (3)
25	Top (16)	Bottom (3)	Middle (1)
26	Top (17)	Bottom (1)	Middle (2)
27	Middle (16)	Bottom (2)	Top (2)


(a) Human faces

(b) Cats


Hair Color

Choose the row in which images are better sorted from black hair (left) to blonde hair (right):

1)



2)



3)



4)



5)



Figure 25: All questions asked in our survey and their associated images. Page 1/6.

Makeup

Choose the row in which images are better sorted from less makeup (left) to more makeup (right):


6)



7)



8)



9)



10)




Figure 25: All questions asked in our survey and their associated images. Page 2/6.

Expression

Choose the row in which images are better sorted from a sad expression (left) to a happy expression (right):

11)



12)



13)



14)



15)





Figure 25: All questions asked in our survey and their associated images. Page 3/6.

Hair Length

Choose the row in which images are better sorted from short hair (left) to long hair (right):


16)




17)



18)



19)




20)



Figure 25: All questions asked in our survey and their associated images. Page 4/6.

Cats Yaw

Choose the row in which images are better sorted according to their left-to-right head angle (yaw):



21)



22)




23)



Figure 25: All questions asked in our survey and their associated images. Page 5/6.

Cats Pitch

Choose the row in which images are better sorted according to their up-to-down head angle (pitch):



24)



25)



26)



27)

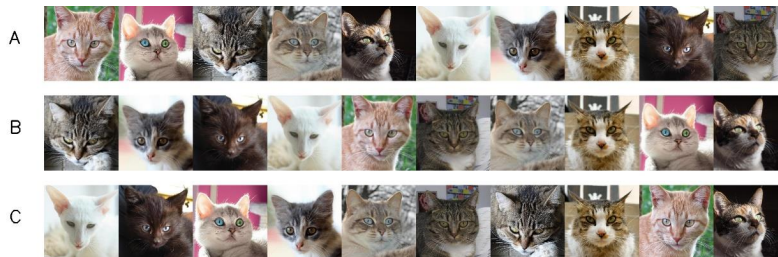


Figure 25: All questions asked in our survey and their associated images. Page 6/6.

D COMPLEMENTING EXPERIMENTS’ DETAILS

We provide additional details about experiments conducted in the main paper.

D.1 ACCURACY OF TRAINED SVMs

In Subsection 4.1 of the paper, we demonstrate that distances in the latent space of the GAN are more semantically meaningful and better behaved than equivalent distances in alternative feature spaces. For this end, we train SVMs in all feature spaces. In Table 3 we report the accuracy of those SVMs on validation sets. As can be seen, the gap in performance for the task of regression cannot be easily explained by the performance of the SVM as a binary classifier. This is further evidence that StyleGAN’s latent space possesses unique properties which make it suitable for the task of regression.

Table 3: Validation accuracy of SVM baseline models operating on pose and age, using the different feature spaces described in Subsection 4.1. As can be seen, most model are decent classifiers.

Feature space	Pose	Age
Ours	0.93	0.82
Pixel	0.87	0.74
Binary-cls	0.91	0.88
GHFeat (Xu et al., 2020)	0.91	0.83
ImageNet (Zagoruyko & Komodakis, 2016)	0.73	0.84
ID (Deng et al., 2019)	0.65	0.85
SwaV (Caron et al., 2020)	0.71	0.88

D.2 WHAT POINTS TO LABEL?

In order to calibrate the latent distances to actual real-world values, a few labeled samples are required. These labeled samples are then used to train a simple linear regression model. In some real-world scenarios, one won’t have pre-defined disjoint sets of labeled and unlabeled samples. Rather, as most datasets form, at first the dataset is simply a collection of unlabeled samples and only later will those be annotated. While our method performs better as it gets more labeled samples, working with a few labeled samples is usually preferred. We thus provide some simple practical suggestions as to what samples one should label. Following these suggestions is increasingly more important as less points are sampled.

First we suggest to invert the unlabeled dataset to the latent space and obtain the distances from the hyperplane corresponding to the semantic latent direction, which results in a distribution of distances. This process does not require any labels. Now, we suggest choosing and labeling a set which is roughly evenly-spaced throughout the center of the distribution. There are two motivations for this sampling strategy, stemming from a single simple principle - sampling points that best represent the distribution. First, samples on the edge of the distribution are more likely to be outliers. Latent outliers may come about when the original image is in itself an outlier in the image distribution. Thus, discarding the noisy edges and sampling from the center of distribution is likely to better represent the dataset. Second, sampling points which are ”close” to each other on the distance axis, is prone to error. The linear relationship between the attribute and distance is modeled by $y = a \cdot d + b + \varepsilon$ where a, b are the function coefficients and ε is an error term. Consider the case of sampling two points with distances d_1, d_2 . When $\Delta d = d_1 - d_2$ is small, it may be the case that $\varepsilon > a \cdot \Delta d + b$. In such case, the noise in the observed attribute may overwhelm any signal due to the modified latent-distance, and a linear model fit to these points will fail to predict the underlying a, b . Sampling a set which is roughly evenly-spaced maximizes the minimal distance between any two points.

We follow these guidelines when conducting all experiments described in the paper. For the center, we consider 95% of the data. For evenly-spaced distances we first observe that for n points sampled from a uniform distribution over $[a, b]$, the minimal distance between a pair of samples is smaller or equal to $\frac{b-a}{n}$. However, choosing such a minimal distance will only allow for, at most, one sampled set. To allow greater flexibility in the choice of samples, we loosen the restriction, and sample points whose minimal allowed difference is $\frac{b-a}{n^{1.3}}$.

D.3 REGRESSION MODEL REGULARIZATION

Our final regression model is a simple linear regression model. However, there is still room to choose regularization. We experiment with our model without regularization, with $L1$ regularization (*i.e.* Lasso), $L2$ regularization (*i.e.* Ridge) or both (*i.e.* ElasticNet). We find that there’s only a slight difference in the few-shot setting and it diminishes as the number of samples increases, as demonstrated in Figure 26. We used ElasticNet regularization in all experiments presented in the paper. We use the default penalty weighting provided by scikit-learn (Pedregosa et al., 2011).




Figure 26: We compare the results of our approach using different types of regularization on the simple linear regression model. As can be observed, only slight difference exist when two calibration points as used, and difference diminishes as more points are sampled.

E LINEARITY ORIGIN HYPOTHESIS

Our method builds on the observation that an approximately linear correlation exists between distances in StyleGAN’s latent space and the magnitude of semantic properties in an image. This observation is largely empirical in nature (see Figures 1 and 8). Nevertheless, we hypothesize on the origin of this property. As datasets are non uniform, the model is tasked with representing values with different densities. One solution would be to “allocate” more space in \mathcal{W} for more frequent values, creating a non-linear space. However, the model already has a mechanism to represent non-uniform densities - through the original, Gaussian distribution of \mathcal{Z} space and it’s non-linear mapping into \mathcal{W} . We thus speculate the model has no incentive to model densities in the \mathcal{W} space itself. Instead, it can model a simpler, linear space, in order to make the generative process simpler. Such a scenario is in line with the unwarping intuition provided in the original StyleGAN paper (Karras et al. (2019)).

F LICENSES AND PRIVACY

The datasets and models used in our work and their respective licenses are outlined in Table 4.

Some of the datasets in use, in particular FFHQ (Karras et al., 2019), CelebA-HQ (Karras et al., 2017) and CACD (Chen et al., 2014), contain personally identifiable data in the form of face images.

We have not reached out to receive consent from the individuals portrayed in the images. However, all three image sets are composed of publicly available celebrity images or faces of individuals crawled from flicker, all of which were uploaded under permissive licenses which allow free use, redistribution, and adaptation for non-commercial purposes. The curators of all sets provide contact details for individuals who wish to have their images removed from the set.

Table 4: Datasets and models used in our work and their respective licenses.

Dataset	Source	License	Model	Source	License
FFHQ	(Karras et al., 2019)	CC BY-NC-SA 4.0*	StyleGAN2	(Karras et al., 2020b)	Nvidia Source Code License-NC
CelebA-HQ	(Karras et al., 2017)	CC BY-NC 4.0	GHFeat	(Xu et al., 2020)	No License
AFHQ	(Choi et al., 2020)	CC BY-NC 4.0	SSV	(Mustikovela et al., 2020)	Nvidia Source Code License-NC
Stanford Cars	(Krause et al., 2013)	ImageNet License	Scikit-Learn	(Pedregosa et al., 2011)	BSD 3-Clause
CompCars	(Yang et al., 2015)	Non-Commercial Research	WHENet	(Zhou & Gregson, 2020)	BSD 3-Clause
CACD	(Chen et al., 2014)	Academic Research	DEX	(Rothe et al., 2015)	No License
PlantVillage	(Hughes et al., 2015)	CC0 1.0	pSp	(Richardson et al., 2020)	MIT License
			e4e	(Tov et al., 2021)	MIT License
			ReStyle	(Alaluf et al., 2021)	MIT License
			InterFaceGAN	(Shen et al., 2020)	MIT License
			SeFa	(Shen & Zhou, 2020)	MIT License
			StyleCLIP	(Patashnik et al., 2021)	MIT License
			CLIP	(Radford et al., 2021)	MIT License
			PoseContrast	(Xiao et al., 2021)	MIT License
			FSA	(Yang et al., 2019)	Apache License V2.0
			StyleGAN2-pytorch	(Seonghyeon, 2020)	MIT License
			StyleGAN-ADA	(Karras et al., 2020a)	Nvidia Source Code License

(a) Datasets

(b) Models