

# ContraFeat: Contrasting Deep Features for Semantic Discovery

Xinqi Zhu, Chang Xu, Dacheng Tao

The University of Sydney, Australia

xzhu7491@uni.sydney.edu.au, c.xu@sydney.edu.au, dacheng.tao@gmail.com

## Abstract

StyleGAN has shown strong potential for disentangled semantic control, thanks to its special design of multi-layer intermediate latent variables. However, existing semantic discovery methods on StyleGAN rely on manual selection of modified latent layers to obtain satisfactory manipulation results, which is tedious and demanding. In this paper, we propose a model that automates this process and achieves state-of-the-art semantic discovery performance. The model consists of an attention-equipped navigator module and losses contrasting deep-feature changes. We propose two model variants, with one contrasting samples in a binary manner, and another one contrasting samples with learned prototype variation patterns. The proposed losses are defined with pre-trained deep features, based on our assumption that the features can implicitly reveal the desired semantic structure including consistency and orthogonality. Additionally, we design two metrics to quantitatively evaluate the performance of semantic discovery methods on FFHQ dataset, and also show that disentangled representations can be derived via a simple training process. Experimentally, our models can obtain state-of-the-art semantic discovery results without relying on latent layer-wise manual selection, and these discovered semantics can be used to manipulate real-world images.

## 1 Introduction

Deep generative models such as GANs (Goodfellow et al. 2014) have shown superior ability in generating photo-realistic images (Karras et al. 2018, 2020; Brock, Donahue, and Simonyan 2019). Besides improving fidelity in generation, many works seek to augment these models with interpretability (Yan et al. 2016; Chen et al. 2016; Higgins et al. 2017; Bau et al. 2019; Shen et al. 2019), which fertilizes applications such as controllable generation (Kulkarni et al. 2015; Lample et al. 2017; Lee et al. 2018; Xing et al. 2019), domain adaptation (Peng et al. 2019; Cao et al. 2018), machine learning fairness (Creager et al. 2019; Locatello et al. 2019a), etc. Yet this goal can be typically achieved via supervised learning (Kingma et al. 2014; Dosovitskiy, Springenberg, and Brox 2014; Kulkarni et al. 2015; Lample et al. 2017; Shen et al. 2019), another recent popular branch of work approaches this target by *semantic discovery*. Instead of training a fresh generator with labels, these semantic

discovery methods (Voynov and Babenko 2020; Härkönen et al. 2020; Wang and Ponce 2021; Shen and Zhou 2021; Yüksel et al. 2021) seek semantically meaningful directions in the latent space of pretrained GANs without supervision. Among popular GANs, StyleGAN has shown great potential for disentangling semantics, thanks to its intermediate  $\mathcal{W}$  space design with a multi-layer structure. This paper focuses on StyleGAN to develop a tailored semantic discovery model in its  $\mathcal{W}$  space.

Although many existing models (Härkönen et al. 2020; Shen and Zhou 2021; Yüksel et al. 2021) have demonstrated ability in discovering semantically meaningful variations in StyleGANs, these methods share a common problem: they rely on manual layer-wise selection in the  $\mathcal{W}$  space. As mentioned in (Härkönen et al. 2020), this layer-wise selection is important for satisfactory manipulation results since entanglement can frequently happen in the discovered directions without this process. Unfortunately, this manual selection process is tedious and demanding because it requires humans to inspect and evaluate a large number of layer-wise modification combinations in the  $\mathcal{W}$  latent space.

There are two major reasons causing this problem in existing models: the first is the lack of a module designed for this layer-selection process; the second is the lack of a criterion that can substitute human intelligence to guide this selection. To solve the first problem, we develop a module to predict layer-wise attentions in the  $\mathcal{W}$  space which simulates the selection of latent layers to apply discovered directions. For the second problem, we adopt a pretrained neural network to serve as the selective guidance. This is inspired by existing observation that the perceptual judgment of deep network features largely agrees with human judgment (Zhang et al. 2018). In this paper, we make the hypothesis that this perceptual agreement also extends to semantic structure, e.g. similarity in *semantic changes* judged by humans is agreed by neural networks, and difference as well. Technically, we develop a contrastive approach to encourage the discovery of semantics that will cause perceptually consistent and orthogonal changes in the pretrained feature maps. We propose two variants of our model, with one contrasting samples in a binary manner, and another one contrasting samples with learned prototype variation patterns. With the proposed models, we can discover more disentangled semantics on StyleGANs than existing methods without relying on

## 2 Related Work

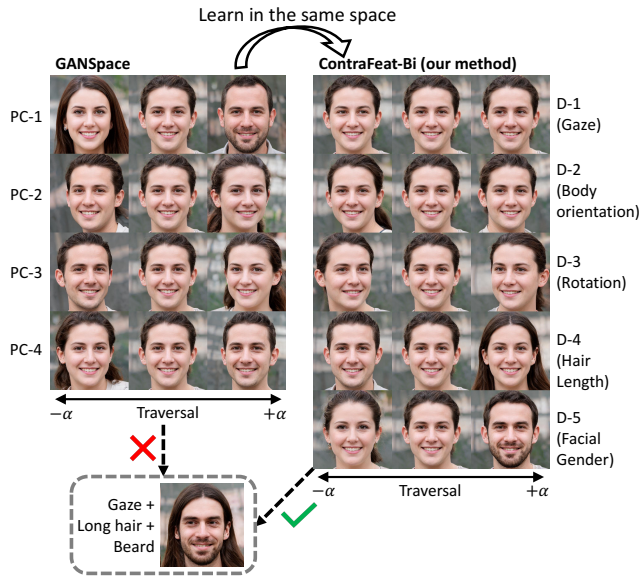


Figure 1: Discovered semantics using Ganspace (Härkönen et al. 2020) and our ContraFeat model in the same  $\mathcal{W}$  subspace (the subspace defined by the top 4 principal components found by Ganspace) of a StyleGAN2 network. No manual  $\mathcal{W}$ -space layer selection is conducted. In this example, our ContraFeat model discovers more disentangled semantics than Ganspace, including the decomposed gender and hair length attributes and the less obvious gaze change. This enables the construction of an out-of-distribution sample *a man with long hair and beard gazing at the left* shown in the bottom-left, which is not possible using Ganspace.

manual layer-wise selection on  $\mathcal{W}$  space. As shown in Fig. 1, both methods (Ganspace (Härkönen et al. 2020) and our model) are applied to the same  $\mathcal{W}$  space, and ours can discover cleaner semantics than the Ganspace model. Note that our model decomposes the hair length and gender attributes into two controllable factors, thus it can manipulate the latent code to generate an out-of-distribution face sample with long-hair and beard (shown in the bottom-left), which is not possible using Ganspace.

Our main contributions are summarized as follows:

- We propose to substitute the manual layer-wise selection procedure in the  $\mathcal{W}$  space of StyleGANs by a learnable attention module.
- We propose to use the deep features in pretrained networks to guide the discovery of semantics, verifying our hypothesis that deep features perceptually agree with human judgment on the semantic structure.
- We design one of our proposed model variants to learn prototype patterns for discovered semantics, avoiding the requirement of sampling multiple changing samples during contrastive learning.
- State-of-the-art semantic discovery results can be achieved with our models, and disentangled representations can be derived.

**Semantic Discovery.** Generative adversarial networks (Goodfellow et al. 2014) have demonstrated state-of-the-art performance (Karras et al. 2018; Brock, Donahue, and Simonyan 2019; Karras, Laine, and Aila 2020; Karras et al. 2020) in generating photo-realistic images. Recently, works revealed that the latent spaces of well-trained GANs can be interpreted with meaningful semantic directions (Bau et al. 2019; Shen et al. 2019). This phenomenon motivated the investigation of semantic properties in the latent spaces (Härkönen et al. 2020; Wang and Ponce 2021), and the discovery of semantic directions (Plumerault, Borgne, and Hudelot 2020; Jahanian, Chai, and Isola 2020; Voynov and Babenko 2020; Spingarn, Banner, and Michaeli 2021; Shen and Zhou 2021; Yüksel et al. 2021). In both (Härkönen et al. 2020) and (Wang and Ponce 2021), the latent space of GANs has been shown highly anisotropic, revealing most semantic variations are encoded in a low-dimension latent subspace. Methods in (Plumerault, Borgne, and Hudelot 2020; Jahanian, Chai, and Isola 2020) successfully learn latent directions to represent predefined transformations on images. Without relying on optimizations, (Härkönen et al. 2020; Spingarn, Banner, and Michaeli 2021; Shen and Zhou 2021) discover semantic attributes by finding the latent directions that cause the most significant changes. Unfortunately, this assumption is too simple and the methods rely on manual adjustment in the latent space to obtain satisfactory results. Based on the idea that different latent directions should correspond to different changes, (Voynov and Babenko 2020) and (Yüksel et al. 2021) adopted classification loss and contrastive loss to encourage the separation of different directions respectively. However, the former requires the training of a deep classifier, and the later is limited to using early features in the generator due to its dense computation in the contrastive loss.

**Disentanglement Learning.** Another popular branch of work to unsupervisedly find semantically meaningful changes is by learning disentangled representations (Higgins et al. 2017; Chen et al. 2016). In this domain, various regularization methods have been proposed based on different assumptions, e.g. statistical independence (Burgess et al. 2018; Kumar, Sattigeri, and Balakrishnan 2018; Kim and Mnih 2018; Chen et al. 2018), informativeness (Chen et al. 2016; Jeon et al. 2021; Zhu, Xu, and Tao 2021b), separability (Lin et al. 2020), and group decomposition (Zhu, Xu, and Tao 2021a). These models usually succeed at the cost of degraded generation quality compared with the normal generative models. In (Locatello et al. 2019b), Locatello et al. proved that unsupervised disentanglement learning is impossible without additional biases. In (Locatello et al. 2020), Locatello et al. showed that guaranteed disentanglement can be achieved with partial factor variation information on some VAE variants (Hosoya 2019; Bouchacourt, Tomioka, and Nowozin 2018). In this work, we show that state-of-the-art disentangled representations can be derived using our semantic discovery model, and the biases in our model include the perceptual contrastive property between semantics, spatial separation, and their hierarchical abstractness difference.

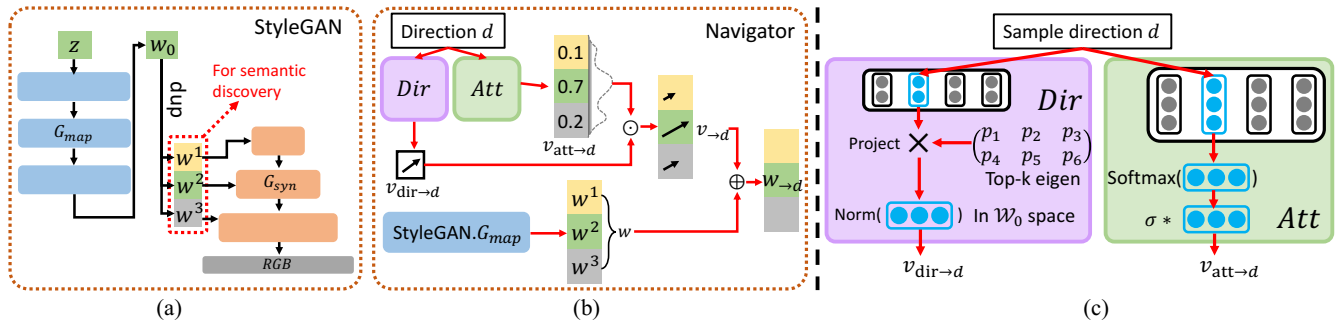


Figure 2: (a) A simplified diagram of StyleGAN. The discovered latent directions are applied on the  $\mathcal{W}$  space after duplication (red dotted box). (b) Overview of our proposed navigator. It consists of two branches. The direction branch ( $Dir$ ) predicts directions  $v_{dir \rightarrow d}$  in the  $\mathcal{W}_0$  space and the attention branch ( $Att$ ) predicts attentions on  $\mathcal{W}$ -space layers  $v_{att \rightarrow d}$ . The combined results  $v_{\rightarrow d}$  are used to modify the latent codes  $w \in \mathcal{W}$ . (c) Detailed illustration of  $Dir$  and  $Att$ .

### 3 Method

In this section, we introduce our semantic discovery model ContraFeat (**C**ontrasting deep **F**eatures) in detail.

#### Navigator

In StyleGANs (a simplified diagram shown in Fig. 2 (a)), a normally sampled latent code  $z$  is fed through a mapping network  $G_{map}$  to obtain an intermediate latent code  $w_0$ . The  $w_0 \in \mathcal{W}_0$  is then broadcast as an extended code  $w \in \mathcal{W}$  via duplication, where  $\mathcal{W}$  is the extended space  $k$  times larger than the original  $\mathcal{W}_0$  space, with  $k$  determined by the depth of the synthesis network ( $G_{syn}$ ). We refer this duplication dimension as *layer* of this  $\mathcal{W}$  space. Thanks to this hierarchical design, disentanglement property emerges in StyleGANs (Karras, Laine, and Aila 2020; Karras et al. 2020). This extended latent space  $\mathcal{W}$  plays an important role in semantic editing, as works (Härkönen et al. 2020; Shen and Zhou 2021) have shown that layer-wise editing in the  $\mathcal{W}$  space achieves more disentangled results. It is also discovered that real-world images can be more accurately inverted into the extended  $\mathcal{W}$  space (Abdal, Qin, and Wonka 2019, 2020) to perform semantic editing.

However, the existing works require manual selection of the modifiable  $\mathcal{W}$ -space layers in order to curate satisfactory editing results. To eliminate this tedious human effort, we propose a navigator network to automate this procedure. Specifically, the module conducts semantic direction prediction (direction branch) and layer-wise selection (attention branch) in the  $\mathcal{W}$  space (depicted in Fig. 2 (c) purple and green boxes respectively). In our design, these two submodules are separated and thus it is possible to replace the direction branch by other existing methods to verify the compatibility of our layer-selection branch with other direction prediction methods.

The direction branch ( $Dir$  purple boxes in Fig. 2) outputs  $m$  latent directions  $\{v_{dir \rightarrow d}, d = 1..m\}$  in the  $\mathcal{W}_0$  space. During training, we randomly sample  $d$ 's to compute losses which will be introduced later. Rather than allowing  $v_{dir \rightarrow d}$  to learn in the whole  $\mathcal{W}_0$  space, we leverage PCA projection to constrain the discovery results in the informative subspace

of  $\mathcal{W}$ , which will be introduced in the next subsection in detail.

To conduct the selection along the  $\mathcal{W}$ -space layer dimension, we build a branch to predict attention weights on this layer dimension (depicted by the green boxes in Fig. 2 (c)). Mathematically,

$$v_{att \rightarrow d} = \sigma * \text{softmax}(\text{att\_logits}_d), \quad d \in 1..m, \quad (1)$$

where the term “att\_logits $_d$ ” denotes the output logits of the attention branch (whose length equal to the  $\mathcal{W}$ -space layers). We use a Gaussian kernel  $\sigma$  convolution to soften the outputs of the softmax function. Without the convolution operation, the modifications tend to be too sharp and are likely to cause artifacts on images. Note that each predicted direction  $v_{dir \rightarrow d}$  corresponds to an attention output  $v_{att \rightarrow d}$ . We compute the product of them to obtain the latent modification, which is then applied to the latent code  $w \in \mathcal{W}$ :

$$v_{\rightarrow d} = v_{dir \rightarrow d} \odot v_{att \rightarrow d}, \quad (2)$$

$$w_{\rightarrow d} = w + v_{\rightarrow d} \quad d \in 1..m, \quad (3)$$

with  $w_{\rightarrow d}$  being the edited code as illustrated in Fig. 2 (b).

#### Informative Subspace in $\mathcal{W}_0$

The  $\mathcal{W}_0$  space of StyleGAN has been shown highly anisotropic (Härkönen et al. 2020; Wang and Ponce 2021), with a small latent subspace contributing to almost all generated varieties (highly informative) while leaving the rest of the space controlling almost nothing. This phenomenon hinders our discovery process as the solution space is sparse. Concerning this issue, we restrict the learning of predicted directions  $v_{dir \rightarrow d}$  in the high-informative subspace so that the navigator can more easily converge to meaningful directions. In practice, this informative subspace is obtained by computing PCA on a large number of samples in the  $\mathcal{W}_0$  space similar to (Härkönen et al. 2020), and using the top- $k$  components (eigenvectors) to define the informative subspace. (see Appendix Sec. A for a detailed introduction).

Let the columns of matrix  $V_{\text{pca}} \in \mathbb{R}^{|\mathcal{W}_0| \times |\mathcal{W}_0|}$  denote the computed principal components and the left-most  $k$  columns

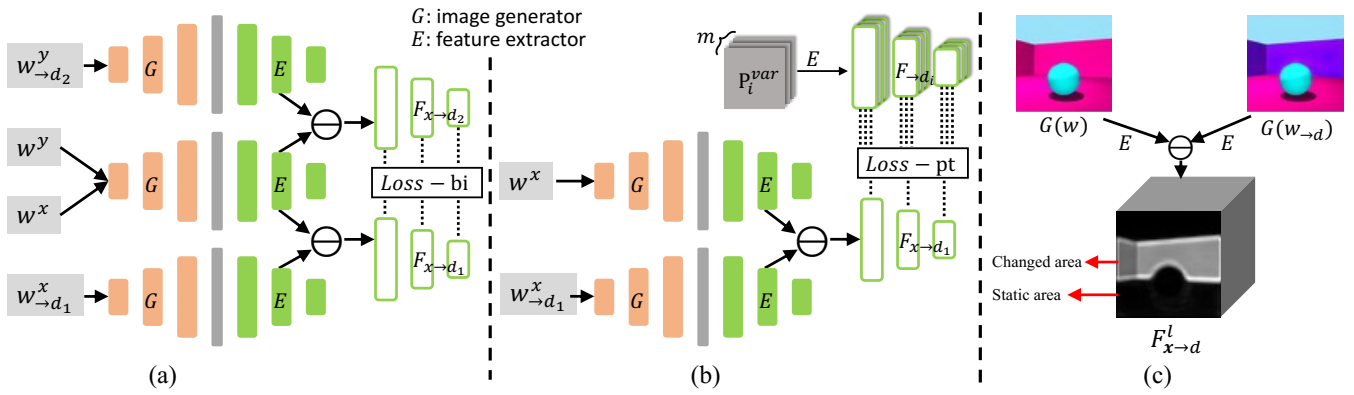


Figure 3: (a). Binary contrasting process. Two directions of change are sampled to contrast with each other. (b). Prototype contrasting process. One direction of change is sampled to contrast with prototype variation patterns. These prototype patterns are learned with gradient decent. (c). A layer of perceptual difference computed by two images differing in a single semantics (wall color). The pixel difference is large for changed areas and small for static areas.

form the basis of the used subspace. Instead of directly learning vectors  $v_{dir \rightarrow d} \in \mathbb{R}^{|\mathcal{W}_0|}$  in the  $\mathcal{W}_0$  space, we learn smaller vectors  $v_{sub \rightarrow d}$  of length  $k$  and project them into the  $\mathcal{W}_0$  space with the top- $k$  PCA basis:

$$v_{dir \rightarrow d} = \text{norm}(V_{\text{pca}}[:, :k] \cdot v_{sub \rightarrow d}), \quad v_{sub \rightarrow d} \in \mathbb{R}^k, \quad (4)$$

where  $\text{norm}$  normalizes the vectors to a constant length. This projection ensures the learned directions are positioned in the informative subspace defined by the top- $k$  principal components. This process is depicted in Fig. 2 (c) purple box.

### Losses

We propose to use pretrained-network features to guide the discovery process. Since deep features show similar judgment to humans (Zhang et al. 2018), we hypothesize this agreement extends to semantic variation structure. In this paper, we consider losses reflecting the simple semantic structure of variation consistency (of the same semantics) and orthogonality (between different semantics).

Given a latent code  $w \in \mathcal{W}$  and its modified version (on direction  $d$ )  $w_{\rightarrow d} = w + v_{\rightarrow d}$ , the corresponding deep-feature changes are defined as:

$$[F_{x \rightarrow d}^l]_{l=1..L} = (E \circ G)(w_{\rightarrow d}) - (E \circ G)(w), \quad 1 \leq d \leq m, \quad (5)$$

where  $G$  is the generator and  $E$  is the feature extractor (e.g. a ConvNet). Usually the extracted features will contain multiple layers ( $l = 1..L$ ). This variation extraction process is illustrated in Fig. 3 (a) (b) (orange and green layers).

**Consistency Loss.** For consistency, we compute two deep-feature changes by applying the same latent direction  $d$  to different latent samples  $x$  and  $y$ , and force them to be similar to each other:

$$l_{\text{cons-avg}}^{\text{bi}} = - \sum_{l=1}^L \left[ \frac{1}{S} \sum_{s=1}^S \text{Sim}^2(F_{x \rightarrow d}^{ls}, F_{y \rightarrow d}^{ls}) \right], \quad (6)$$

where  $s$  indexes the spatial positions of the feature maps, and  $\text{Sim}$  is the cosine similarity function. Unfortunately, this design has a flaw: it aggregates scores on all positions evenly

but only varied areas are meaningful to the consistency measurement. An example of this uneven case is shown in Fig. 3 (c) where only the wall color has varied and the rest areas are static. Since measuring variation similarity on static areas is meaningless, we instead adopt an L2-norm mask of the feature-map difference to implement the weighted aggregation:

$$l_{\text{cons}}^{\text{bi}} = - \sum_{l=1}^L \left[ \frac{1}{Q} \sum_{s=1}^S q(s) \text{Sim}^2(F_{x \rightarrow d}^{ls}, F_{y \rightarrow d}^{ls}) \right], \quad (7)$$

$$q(s) = \frac{\text{nm}(F_{x \rightarrow d}^s)}{\text{nm}_{\max}(F_{x \rightarrow d}^s)} \cdot \frac{\text{nm}(F_{y \rightarrow d}^s)}{\text{nm}_{\max}(F_{y \rightarrow d}^s)}, \quad Q = \sum_{s=1}^S q(s), \quad (8)$$

where  $\text{nm}(F_{x \rightarrow d}^s)$  computes the L2-norm of the activations on  $F_{x \rightarrow d}$  at position  $s$ . This weighted aggregation forces the measurement to focus on the overlapped changes in both inputs. The division by  $Q$  normalizes the weighted aggregation so that the loss cannot favor large changed areas. Empirically this version of consistency measurement leads to better discovery results. We compare different designs in experiment.

Eq. 8 compares variations in a binary manner, where we need to apply the same change to two different codes (illustrated in Fig. 3 (a)), which is not very efficient due to the doubled computation in generator. To solve this issue, we implement another contrastive learning variant by learning a set of prototype patterns:

$$l_{\text{cons}}^{\text{pt}} = - \sum_{l=1}^L \left[ \frac{1}{Q} \sum_{s=1}^S q(s) \text{Sim}^2(F_{x \rightarrow d}^{ls}, F_{* \rightarrow d}^{ls}) \right], \quad (9)$$

$$[F_{* \rightarrow d}^l]_{l=1..L} = E(\tanh(P_d^{var})), \quad (10)$$

where  $P_d^{var} \in \mathbb{R}^{h \times w \times c}$  denotes the prototype pattern of semantic  $d$ , learned along with the model. These prototype patterns serve as hallucinated pictures of the discovered semantic changes in the RGB space (see visualizations in Fig. 5), which are supposed to cause similar deep-feature changes as by actual image pairs.




Figure 4: Fixing discovered  $\mathcal{W}_0$ -space directions predicted by Ganspace (Härkönen et al. 2020), we compare the difference introduced by our proposed attention branch. We can see by incorporating the attention module, the interpretability of each direction is improved. For instance, in the first direction, the facial gender features and skin color shift are removed, resulting in a specific control on hair length. (Bottom-right) Predicted  $W$ -space layer attentions for some semantics.

Models	$S_{\text{disen}}$	$N_{\text{discov}}$
Gsp_dir-non-non	0.304	9
Gsp_dir-att-bi	$0.320 \pm 0.014$	$10.6 \pm 1.02$
Gsp_dir-att-pt	<b><math>0.327 \pm 0.015</math></b>	<b><math>11.8 \pm 0.75</math></b>
Sefa_dir-non-non	0.290	9
Sefa_dir-att-bi	<b><math>0.328 \pm 0.023</math></b>	<b><math>13.8 \pm 0.75</math></b>
Sefa_dir-att-pt	$0.321 \pm 0.016$	$13.0 \pm 0.63$

Table 1: Comparing the semantic discovery performance with and without our attention module on baselines Ganspace (Härkönen et al. 2020), and Sefa (Shen and Zhou 2021). The shown model names follow the pattern: dir\_module - att\_module - loss\_type.

**Orthogonality Loss.** Similarly, we define the orthogonality losses (both binary and prototype implementations), which encourage feature variations caused by different latent changes to be orthogonal:

$$l_{\text{orth}}^{\text{bi}} = \sum_{l=1}^L \left[ \frac{1}{Q} \sum_{s=1}^S q(s) \text{Sim}^2(F_{\mathbf{x} \rightarrow d}^{ls}, F_{\mathbf{y} \rightarrow d'}^{ls}) \right], \quad (11)$$

$$l_{\text{orth}}^{\text{pt}} = \frac{1}{m-1} \sum_{d' \neq d} \sum_{l=1}^L \left[ \frac{1}{Q} \sum_{s=1}^S q(s) \text{Sim}^2(F_{\mathbf{x} \rightarrow d}^{ls}, F_{\mathbf{x} \rightarrow d'}^{ls}) \right], \quad (12)$$

where  $d, d'$  denote different directions predicted by our navigator.

**Diversity Regularization.** We add a small regularization term to encourage the navigator to proposed different directions in the latent space. This term practically increases the diversity of discovered concepts:

$$l_{\text{div}} = \sum_{i,j} \text{Sim}^2(\mathbf{v}_{\text{dir} \rightarrow i}, \mathbf{v}_{\text{dir} \rightarrow j}), \quad i, j \in 1..m, i \neq j. \quad (13)$$

In summary, our overall losses (both binary and prototype) are:

$$l^{\text{bi}} = l_{\text{cons}}^{\text{bi}} + l_{\text{orth}}^{\text{bi}} + \lambda l_{\text{div}} \quad \text{and} \quad l^{\text{pt}} = l_{\text{cons}}^{\text{pt}} + l_{\text{orth}}^{\text{pt}} + \lambda l_{\text{div}}. \quad (14)$$

## 4 Experiments

### Evaluation Metrics

We design two quantitative metrics on the FFHQ dataset (Karras, Laine, and Aila 2020) to evaluate the semantic discovery performance of models. Considering there are no

labeled attributes on this dataset, we leverage a pretrained state-of-the-art disentanglement model (Zhu, Xu, and Tao 2021b) to construct an attribute predictor. Specifically, for images  $X$  generated by the disentangled generator, we train a predictor  $P(X)$  to reconstruct the disentangled latent codes  $z \in R^p$ , with each dimension representing a score for an attribute (total of  $p$  attributes).

To evaluate a semantic discovery model, we iterate through all the predicted directions of the model, apply these modifications to a set of  $N$  samples, and compute their mean absolute difference on the scores predicted by the pretrained predictor  $P$ . Mathematically, for  $d = 1..m$ :

$$A[d, :] = \frac{A'[d, :]}{\max(A'[d, :])}, \quad (15)$$

$$A'[d, :] = \frac{1}{N} \sum_{i=1}^N |(P \circ G)(\mathbf{w}_i + \mathbf{v}_{\rightarrow d}) - (P \circ G)(\mathbf{w}_i)|, \quad (16)$$

where each entry of the matrix  $A \in R^{m \times p}$  represents how an attribute score ( $i = 1..p$ ) changes on average when images vary along a discovered semantic direction ( $d = 1..m$ ). Note that this matrix is normalized so that the maximum value of each row is 1. Then we define two evaluation quantities:

$$S_{\text{disen}} = \frac{1}{m} \sum_{d=1}^m \max_{1\text{st}}(A[d, :]) - \max_{2\text{nd}}(A[d, :]), \quad (17)$$

$$N_{\text{discov}} = \sum_{i=1}^p \mathbf{1}_{\{x | 1 \in x\}}(A[:, i]), \quad (18)$$

where  $\max_{1\text{st}}(A[d, :])$  extracts the largest value in  $A[d, :]$ , and  $\max_{2\text{nd}}(A[d, :])$  extracts the second largest value. The  $S_{\text{disen}}$  evaluates how clean a discovered semantics is because an ideal discovery should only cause a single ground-truth attribute to change. The  $N_{\text{discov}}$  quantity counts the number of semantics discovered by the model. The term  $\mathbf{1}_{\{x | 1 \in x\}}(A[:, i])$  is 1 if the attribute  $i$  is at least once predicted as the maximum value in a row (indicating the attribute  $i$  has been discovered by the model). Some additional information of evaluation metrics can be found in Appendix C.

### Effectiveness of the Attention Branch

In Table 1, we quantitatively compare the impact made by our attention branch when applied to Ganspace (Gsp) (Härkönen et al. 2020) and Sefa (Shen and Zhou 2021) baselines. The results are obtained by training a proposed attention module on the top-30 directions predicted by either base

Models	$S_{\text{disen}}$	$N_{\text{discov}}$
Bi-30	$0.313 \pm 0.015$	$13.6 \pm 0.80$
Bi-100	<b><math>0.357 \pm 0.017</math></b>	$14.2 \pm 0.85$
Bi-300	$0.340 \pm 0.012$	<b><math>14.4 \pm 0.80</math></b>
Bi-500	$0.337 \pm 0.014$	$12.0 \pm 0.89$
Pt-30	$0.314 \pm 0.018$	<b><math>14.6 \pm 0.49</math></b>
Pt-100	<b><math>0.344 \pm 0.022</math></b>	$14.2 \pm 0.75$
Pt-300	$0.336 \pm 0.013$	$13.6 \pm 0.49$
Pt-500	$0.332 \pm 0.015$	$11.8 \pm 0.75$

Table 2: Comparing the discovery performance of our models learned in PCA subspaces of different sizes.

models in both binary (bi) and prototype-pattern (pt) implementations. We can see the attention branch improves the baselines on both metrics, especially  $N_{\text{discov}}$ . This is because the base models predict many directions representing overlapped semantics, leading to low total counted number of discovered attributes ( $N_{\text{discov}}$ ), while our attention module removes some of these overlapped semantics, leading to the emerge of other smaller but diverse attributes. In Fig. 4, we qualitatively show the difference before and after the use of the attention branch on the Ganspace predicted directions. We can see though the Ganspace can propose directions to control meaningful image changes, many of them are still mixture of multiple semantics. In contrast, our learned attentions help to remove some correlated changes like the facial gender features, background changes, and hair style changes, leading to cleaner semantic controls.

In Fig. 4 bottom-right, we show the predicted attentions for some semantic concepts discovered on FFHQ. We can see the that high-level changes (e.g. rotation) correspond to  $\mathcal{W}$ -space top layers, facial attribute changes (e.g. smile, hair-color) correspond to middle layers, and changes like ambient hue correspond to low layers, which is similar to results found in (Karras, Laine, and Aila 2020).

### Informative Subspace

As introduced in Sec. 3, there exists an informative subspace in the  $\mathcal{W}_0$  space that accounts for most of the generation variety, and we can restrict the discovery process in this dense subspace to make the discovery easier. In Table 2, we compare the discovery results obtained in subspaces of different sizes. We can see when we restrict the subspace to be very small (30-dim), the disentanglement score  $S_{\text{disen}}$  becomes the lowest in both model variants, while the number of discovered attributes  $N_{\text{discov}}$  becomes highest. This is because the low-dim subspace is content-dense, thus the model can more easily find meaningful changes, leading to higher number in  $N_{\text{discov}}$ . However, the semantics discovered in this small subspace may be incomplete as some of them are clipped and may be embedded beyond this subspace, resulting in partial discovery or artifacts, leading to low  $S_{\text{disen}}$ . On the other hand, using a large subspace (500-dim) leads to lower  $N_{\text{discov}}$ , meaning it becomes harder to converge to meaningful directions as a large part of the space controls

Model	$S_{\text{disen}}$	$N_{\text{discov}}$
Recog	$0.249 \pm 0.023$	$4.8 \pm 0.75$
Recog-Att	$0.312 \pm 0.015$	$8.6 \pm 0.80$
LatCLR	$0.267 \pm 0.014$	$9.6 \pm 1.02$
LatCLR-Att	$0.304 \pm 0.020$	$10.2 \pm 1.17$
Contra-Bi	<b><math>0.357 \pm 0.017</math></b>	<b><math>14.2 \pm 0.85</math></b>
Contra-Pt	<b><math>0.344 \pm 0.022</math></b>	<b><math>14.2 \pm 0.75</math></b>

Table 3: Semantic discovery comparison between different models with and without the proposed attention module on FFHQ dataset (Karras, Laine, and Aila 2020).

null changes. In summary, we find that using a subspace of a middle size (100 or 300-dim) leads to more balanced discovery results.

### Prototype Variation Patterns

In Fig. 5, we present the prototype variation patterns learned by our Contra-Pt model variant. We can see the model successfully learns to characterize many interpretable patterns including some subtle changes like clothes color, eye opening, and glasses. It is also interesting to see that the model spontaneously decomposes the semantics spatially so that they are less like to interfere with each other. Learning such hallucinated patterns in the RGB space opens the possibility for conditional discovery and fine-grained manipulation, as it is possible to constrain the patterns by some predefined semantic/attribute segmentations. We leave this interesting extension for future work.

### Compared to Other Discovery Models

In Table 4 left, we quantitatively compare the proposed models with other existing models including recognition (Recog) (Voynov and Babenko 2020) and LatentCLR (LatCLR) (Yüksel et al. 2021). We can see: (1) our proposed contrastive loss performs best among all models; (2) existing models can be significantly boosted by incorporating our proposed  $\mathcal{W}$ -space attention module. Besides the performance gain, it is also worth noting that (1) the Recog model requires to train an extra deep predictor for the discovery process which takes more training steps to converge while our models only need to train the navigator (and the variation patterns in the PT variant); (2) the LatCLR model is designed to contrast every direction change with each other (an outer product) in very loss computation, which is highly memory-consuming. It prevents the model from using features beyond the early few layers of the generator for its loss computation, limiting the resulted semantic discovery performance. In Table. 4 right, we also compare these losses on disentangled representation learning on 3DShapes dataset (Kim and Mnih 2018). We can see our models still outperform the baselines by an obvious margin.

### Learning Disentangled Representations

With semantics discovered, we can then construct disentangled representations by learning an encoder that aligns the




Figure 5: Prototype variation patterns learned with our proposed Contra-Pt variant. These patterns learn to characterize the changes that are related to the corresponding discovered semantics.

Model	FVM	MIG	DCI
$\beta$ -VAE	90.2 $\pm$ 8.1	39.5 $\pm$ 8.4	67.4 $\pm$ 3.7
FacVAE	91.4 $\pm$ 5.0	39.2 $\pm$ 8.9	<b>76.3</b> $\pm$ 8.0
Recog-Att	86.7 $\pm$ 5.2	29.1 $\pm$ 10.2	49.4 $\pm$ 8.3
LatCLR-Att	92.9 $\pm$ 7.0	40.4 $\pm$ 5.3	70.5 $\pm$ 12.0
Contra-Bi	<b>93.4</b> $\pm$ 6.2	<b>43.8</b> $\pm$ 7.3	72.0 $\pm$ 5.4
Contra-Pt	<b>95.3</b> $\pm$ 4.3	<b>47.2</b> $\pm$ 6.5	<b>74.5</b> $\pm$ 7.0

Table 4: Disentangled representation learning comparison on 3DShapes dataset (Kim and Mnih 2018).

semantic directions with axes in a latent space. This can be achieved by training a model on a dataset composed of paired images, with each data pair differing in a single semantic direction. This dataset can be constructed as follows:

$$\{(G(\mathbf{w}), G(\mathbf{w} + \mathbf{v}_d)) \mid \mathbf{w} \sim p(\mathbf{w}), d = 1..m\}. \quad (19)$$

Then we can train an off-the-shelf group-based VAE model (Hosoya 2019; Bouchacourt, Tomioka, and Nowozin 2018; Locatello et al. 2020) to obtain disentangled representations. The implementation of the group-based VAE is shown in the Appendix Sec. B. With disentangled representations, we can use quantitative metrics (Kim and Mnih 2018; Eastwood and Williams 2018; Chen et al. 2018) to indirectly evaluate the quality of discovered semantic directions.

We show the state-of-the-art comparison on 3DShapes in Table 4. We can see the derived representations by our ContraFeat models match or outperform the existing state-of-the-art disentanglement methods, showing this new pipeline (discovering semantics + disentangling) is an effective unsupervised disentangled representation learning approach. Qualitative results are in Appendix Sec. J.

### Real-image Editing and More Qualitative Results

We can achieve real-image editing by projecting real images into the  $\mathcal{W}$  latent space of StyleGAN. We adopt the projection method from (Karras et al. 2020), with the only modification of learning the projected code in the extended  $\mathcal{W}$  space rather than the originally-used  $\mathcal{W}_0$  space. After projection, we can apply our discovered semantic directions to the projected code to edit the image. In Fig. 6 (left), we show some editing results on real-world images (more results are shown in the Appendix Sec. H). Here we see that we can conduct interesting editing like fatness and nose length adjustment, which are rarely-seen in existing other semantic discovery methods.

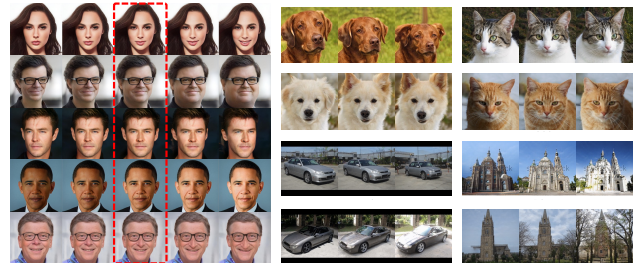


Figure 6: (Left) Some real-image editing results. The red box shows the reconstructed original images. (Right) Some qualitative discovery results on different datasets.

In Fig. 6 right, we show some discovered semantics on other datasets. We can see our model can discover specific semantics for different datasets like the *ear pose* for Dogs dataset and *trees* for Church dataset, indicating our models can work adaptively on these datasets. More discovery results can be found in the Appendix Sec. I.

## 5 Conclusion

Concerning the tedious procedure of manual selection on  $\mathcal{W}$ -space layers in existing semantic discovery methods, we introduced an attention-equipped navigator module to solve this problem. In order to encourage more effective semantic discovery, we additionally proposed to restrict the discovery process in the latent subspace of  $\mathcal{W}$  defined by the top- $k$  principal components. Leveraging the high perceptual capability of deep-network features, we adopted a pretrained network as a cheap semantic sensor to guide the discovery process in a contrastive learning manner by encouraging consistency and orthogonality in image variations. Two model variants were introduced with one contrasting variation samples in a binary way and another contrasting samples with learned prototype variation patterns. To quantitatively evaluate semantic discovery models without ground-truth labels, we designed two metrics on FFHQ based on a pretrained disentanglement model. Experimentally, we showed that the two proposed models are both effective and can achieve state-of-the-art semantic discovery performance. Additionally, we showed that disentangled representations can be derived from our discovered semantics, and real-image editing can also be achieved based on image projection techniques.

## Acknowledgements

This work was supported by the Australian Research Council under projects FL-170100117, IC-190100031, LE-200100049, DP210101859, the University of Sydney Research Accelerator (SOAR) Prize, and the University of Sydney Faculty of Engineering PhD Completion Award.

## References

- Abdal, R.; Qin, Y.; and Wonka, P. 2019. Image2StyleGAN: How to Embed Images Into the StyleGAN Latent Space? In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 4431–4440.
- Abdal, R.; Qin, Y.; and Wonka, P. 2020. Image2StyleGAN++: How to Edit the Embedded Images? In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8293–8302.
- Bau, D.; Zhu, J.-Y.; Strobel, H.; Zhou, B.; Tenenbaum, J. B.; Freeman, W. T.; and Torralba, A. 2019. GAN Dissection: Visualizing and Understanding Generative Adversarial Networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Bouchacourt, D.; Tomioka, R.; and Nowozin, S. 2018. Multi-Level Variational Autoencoder: Learning Disentangled Representations From Grouped Observations. In McIlraith, S. A.; and Weinberger, K. Q., eds., *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 2095–2102. AAAI Press.
- Brock, A.; Donahue, J.; and Simonyan, K. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *ICLR*.
- Burgess, C. P.; Higgins, I.; Pal, A.; Matthey, L.; Watters, N.; Desjardins, G.; and Lerchner, A. 2018. Understanding disentangling in beta-VAE. *ArXiv*, abs/1804.03599.
- Cao, J.; Katzir, O.; Jiang, P.; Lischinski, D.; Cohen-Or, D.; Tu, C.; and Li, Y. 2018. DiDA: Disentangled Synthesis for Domain Adaptation. *ArXiv*, abs/1805.08019.
- Chen, R. T. Q.; Li, X.; Grosse, R.; and Duvenaud, D. 2018. Isolating Sources of Disentanglement in Variational Autoencoders. In *Advances in Neural Information Processing Systems*.
- Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. In *NIPS*.
- Creager, E.; Madras, D.; Jacobsen, J.-H.; Weis, M. A.; Swersky, K.; Pitassi, T.; and Zemel, R. S. 2019. Flexibly Fair Representation Learning by Disentanglement. In *ICML*.
- Dosovitskiy, A.; Springenberg, J. T.; and Brox, T. 2014. Learning to generate chairs with convolutional neural networks. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1538–1546.
- Eastwood, C.; and Williams, C. K. I. 2018. A Framework for the Quantitative Evaluation of Disentangled Representations. In *ICLR*.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A. C.; and Bengio, Y. 2014. Generative Adversarial Networks. In *NIPS*.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M. M.; Mohamed, S.; and Lerchner, A. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *ICLR*.
- Hosoya, H. 2019. Group-based Learning of Disentangled Representations with Generalizability for Novel Contents. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 2506–2513. International Joint Conferences on Artificial Intelligence Organization.
- Härkönen, E.; Hertzmann, A.; Lehtinen, J.; and Paris, S. 2020. GANSpace: Discovering Interpretable GAN Controls. In *Proc. NeurIPS*.
- Jahani, A.; Chai, L.; and Isola, P. 2020. On the “steerability” of generative adversarial networks. In *International Conference on Learning Representations*.
- Jeon, I.; Lee, W.; Pyeon, M.; and Kim, G. 2021. IB-GAN: Disentangled Representation Learning with Information Bottleneck Generative Adversarial Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(9): 7926–7934.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *ICLR*.
- Karras, T.; Laine, S.; and Aila, T. 2020. A Style-Based Generator Architecture for Generative Adversarial Networks. *IEEE transactions on pattern analysis and machine intelligence*.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and Improving the Image Quality of StyleGAN. *CVPR*.
- Kim, H.; and Mnih, A. 2018. Disentangling by Factorising. In *ICML*.
- Kingma, D. P.; Mohamed, S.; Rezende, D. J.; and Welling, M. 2014. Semi-supervised Learning with Deep Generative Models. In *NIPS*.
- Kulkarni, T. D.; Whitney, W. F.; Kohli, P.; and Tenenbaum, J. B. 2015. Deep Convolutional Inverse Graphics Network. In *NIPS*.
- Kumar, A.; Sattigeri, P.; and Balakrishnan, A. 2018. Variational Inference of Disentangled Latent Concepts from Unlabeled Observations. In *ICLR*.
- Lample, G.; Zeghidour, N.; Usunier, N.; Bordes, A.; Denoyer, L.; and Ranzato, M. 2017. Fader Networks: Manipulating Images by Sliding Attributes. In *NIPS*, volume abs/1706.00409.
- Lee, H.-Y.; Tseng, H.-Y.; Huang, J.-B.; Singh, M. K.; and Yang, M.-H. 2018. Diverse Image-to-Image Translation via Disentangled Representations. *ECCV*, abs/1808.00948.
- Lin, Z.; Thekumparampil, K. K.; Fanti, G.; and Oh, S. 2020. InfoGAN-CR: Disentangling Generative Adversarial Networks with Contrastive Regularizers. *ICML*.

Locatello, F.; Abbati, G.; Rainforth, T.; Bauer, S.; Schölkopf, B.; and Bachem, O. 2019a. On the Fairness of Disentangled Representations. In *NeurIPS*.

Locatello, F.; Bauer, S.; Lucic, M.; Rätsch, G.; Gelly, S.; Schölkopf, B.; and Bachem, O. 2019b. Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations. In *ICML*.

Locatello, F.; Poole, B.; Raetsch, G.; Schölkopf, B.; Bachem, O.; and Tschannen, M. 2020. Weakly-Supervised Disentanglement Without Compromises. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 6348–6359. PMLR.

Peng, X.; Huang, Z.; Sun, X.; and Saenko, K. 2019. Domain Agnostic Learning with Disentangled Representations. In *ICML*.

Plumerault, A.; Borgne, H. L.; and Hudelot, C. 2020. Controlling generative models with continuous factors of variations. In *International Conference on Learning Representations*.

Shen, Y.; Gu, J.; Tang, X.; and Zhou, B. 2019. Interpreting the Latent Space of GANs for Semantic Face Editing. *ArXiv*, abs/1907.10786.

Shen, Y.; and Zhou, B. 2021. Closed-Form Factorization of Latent Semantics in GANs. In *CVPR*.

Spingarn, N.; Banner, R.; and Michaeli, T. 2021. {GAN} "Steerability" without optimization. In *International Conference on Learning Representations*.

Voynov, A.; and Babenko, A. 2020. Unsupervised discovery of interpretable directions in the gan latent space. In *International Conference on Machine Learning*, 9786–9796. PMLR.

Wang, B.; and Ponce, C. R. 2021. A Geometric Analysis of Deep Generative Image Models and Its Applications. In *International Conference on Learning Representations*.

Xing, X.; Han, T.; Gao, R.; Zhu, S.-C.; and Wu, Y. N. 2019. Unsupervised Disentangling of Appearance and Geometry by Deformable Generator Network. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10346–10355.

Yan, X.; Yang, J.; Sohn, K.; and Lee, H. 2016. Attribute2Image: Conditional Image Generation from Visual Attributes. In *ECCV*.

Yüksel, O. K.; Simsar, E.; Er, E. G.; and Yanardag, P. 2021. LatentCLR: A Contrastive Learning Approach for Unsupervised Discovery of Interpretable Directions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 14263–14272.

Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.

Zhu, X.; Xu, C.; and Tao, D. 2021a. Commutative Lie Group VAE for Disentanglement Learning. In *ICML*.

Zhu, X.; Xu, C.; and Tao, D. 2021b. Where and What? Examining Interpretable Disentangled Representations. In *CVPR*.