

Constructive Assimilation: Boosting Contrastive Learning Performance through View Generation Strategies

Ligong Han^{1,2*} Seungwook Han² Shivchander Sudalairaj² Charlotte Loh³
 Rumen Dangovski³ Fei Deng¹ Pulkit Agrawal⁴ Dimitris Metaxas¹
 Leonid Karlinsky² Tsui-Wei Weng⁵ Akash Srivastava²

¹Rutgers University ²MIT-IBM Watson AI Lab ³MIT EECS ⁴MIT CSAIL ⁵UCSD

Abstract

Transformations based on domain expertise (expert transformations), such as random-resized-crop and color-jitter, have proven critical to the success of contrastive learning techniques such as SimCLR. Recently, several attempts have been made to replace such domain-specific, human-designed transformations with generated views that are learned. However for imagery data, so far none of these view generation methods has been able to outperform expert transformations. In this work, we tackle a different question: instead of replacing expert transformations with generated views, can we constructively assimilate generated views with expert transformations? We answer this question in the affirmative and propose a view generation method and a simple, effective assimilation method that together improve the state-of-the-art by up to $\approx 3.6\%$ on three different datasets. Importantly, we conduct a detailed empirical study that systematically analyzes a range of view generation and assimilation methods and provides a holistic picture of the efficacy of learned views in contrastive representation learning.

1. Introduction

Contrastive learning (CL) has become one of the most powerful tools for self-supervised learning. Most contrastive methods are trained using instance discrimination: pulling positive views (generated from the same image) close in the learned representation space, while pushing negative views (generated from other images) away [2, 3, 5–8, 12, 15, 16, 21, 28, 32, 34, 36]. Intuitively, the quality of representations learned by CL is highly dependent on the mechanisms for generating these views, which are commonly compositions of a set of handcrafted transformations designed by human experts. The choice of the transformations and the specific design of their composition

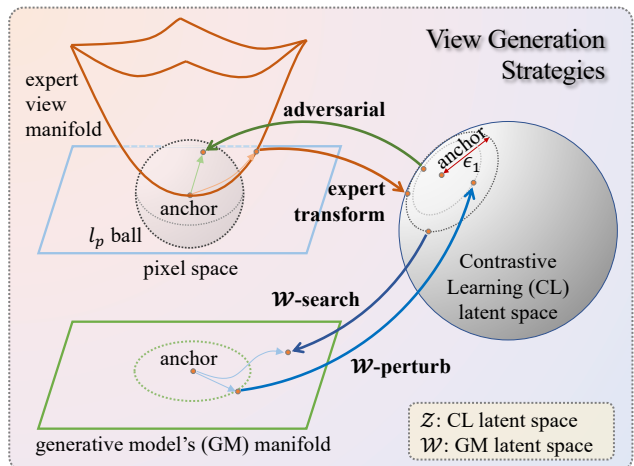


Figure 1. Illustrative visualization of the view generation strategies. The convex shaped surface represents the expert view manifold, and the upper blue and lower green planes represent the pixel space and the GAN’s manifold, respectively. The hypersphere represents the learned representation space of a CL encoder. \mathcal{W} -search searches in the CL model’s latent space (\mathcal{Z} -space) using the loss in Eq. (2) by optimizing w . \mathcal{W} -perturb directly perturbs the GAN’s latent space (\mathcal{W} -space) with Gaussian noise. **Adversarial** (or ViewMaker [27]) performs adversarial training and constrains the generated view in an l_p ball around the anchor in pixel space. **Expert transform** generates views by applying a set of transformations commonly used in CL.

relies on domain expertise built by years of CL research, which mostly focused on general imagery, commonly captured with consumer cameras and collected from the web. However, this domain expertise may not necessarily hold for new unseen visual domains that could be encountered by CL practitioners. Therefore, a natural question arises: *How can we go beyond these expert-designed views?*

A recent theoretical analysis of CL [38] states that CL inverts the true data generative process. This motivates us to explore the use of a generative model (a GAN in this

*Work done during an internship at MIT-IBM Watson AI Lab.

paper) as a view generator for training a contrastive encoder. There are previous works that study view generation through the lens of *generative models*. Viewmaker [27] employs an image-to-image network and generates views via adversarial training. The generated perturbation is bounded in a l_p ball in the pixel space, limiting its applicability to CL that commonly requires strong augmentation for best results. Recently, [17] defines views via a black-box generative model by perturbing slightly in the generative model’s latent space without explicitly controlling the effect of this perturbation in the contrastive encoder latent space. According to analysis in [38], in the ideal case, the generative and the contrastive encoder’s latent spaces should be related by an orthonormal transformation, making transformations in the GAN’s latent space and in the contrastive encoder’s latent space equivalent up to rotation. Moreover, they empirically show that such an equivalence holds to some extent when the theoretical conditions are not satisfied.

2. Background and Related Work

SimCLR [5, 6] is one of the first and most established CL baselines. Let $x \in \mathcal{X}$ denote an image in a mini-batch $\mathcal{B} = \{x_i\}_{i=1}^N$. Further, let $f : \mathcal{X} \mapsto \mathcal{W}$ be a representation encoder, parameterized as a deep neural network. Here $\mathcal{X} \subseteq \mathbb{R}^D, \mathcal{W} \subseteq \mathbb{R}^K$ such that $K < D$. Following the notations from [20], we use $i \in \mathcal{I} \equiv \{1, \dots, 2N\}$ to denote the index of an arbitrary batch (augmented using *expert transformations* as explained below), where $j(i)$ is the index of the other augmented sample originating from the same data sample. $\mathcal{A}(i) := \mathcal{I} \setminus \{i\}$ is the complement of i . The contrastive loss function for SimCLR is the InfoNCE loss, which is

$$L_{\text{simclr}} = - \sum_{i \in \mathcal{I}} \log \frac{\exp(z_i \cdot z_{j(i)}/\tau)}{\sum_{a \in \mathcal{A}(i)} \exp(z_i \cdot z_a/\tau)} \quad (1)$$

where, $z = f(x)$ and τ is the temperature parameter.

Views In contrastive learning, a stochastic data augmentation module applies two different sets of transformations to generate two correlated views of the same data point, x_i and $x_{j(i)}$. The two sets of transformations, t_1 and t_2 , are sampled from the same family of transformations \mathcal{T} , such as cropping and resizing, horizontal flipping, color distortion, etc. The images x_i and $x_{j(i)}$ are the positive views, whereas all other views generated from other samples are the negative views. As studied by many recent works [24, 29, 35], \mathcal{T} defines the invariances that the model learns. Therefore, human domain knowledge and extensive years of research have been invested to search for the optimal composition of augmentations and their corresponding parameters to optimize CL methods for a given data. We refer to the optimal

set of transformations in SimCLR as *expert transformations* and views created by them as *expert views*.

View Generation Previous works have explored the idea of generating views for SSL. ADIOS [26] and Viewmaker [27] both study adversarial methods. ADIOS learns a masking function and an image encoder performing a min-max optimization on the same objective function in the masked image model framework. Alternatively, Viewmaker learns a bounded perturbation directly in the pixel space by employing the min-max adversarial training. Relatedly, [17] leverages a pre-trained GAN model to generate the data samples and to replace the entire training dataset like ImageNet [10] with a generated one.

View Assimilation Previous works like CMC [29] and DINO [3] study the case of having more than two views in contrastive learning. Both works treat this as a special case in which an arbitrary batch of views has been expanded and demonstrate gains from including additional views. However, empirically we observed that simply augmenting SimCLR with an additional set of positive views degrades its performance. Instead, we explore and propose new methods that can effectively integrate additional views and boost performance.

3. View Generation

3.1. \mathcal{W} -search: perturbation in CL’s latent space

Contrastive learning under the InfoNCE loss leads to congregation of the normalized representations $z/\|z\|$ of the positive views around their respective anchors on the hypersphere that they reside on. Expectedly, the distribution of the distances of positive views around their anchor is very similar for different anchors. Fig. 1 illustrates this for the case of $z \in \mathbb{R}^3$. Consequently, a straightforward strategy is to generate more positive views from other points in \mathcal{Z} that are the same distance away from the nearest anchor. The main challenge, however, is how to find the corresponding point in the \mathcal{X} space to realize the view?

We propose a simple technique to resolve this challenge by leveraging a pre-trained GAN generator g [19]. This allows us to pose the view generation as an optimization problem, *i.e.* find a w^* that minimizes the following loss function. Assuming we are generating n views simultaneously,

$$\{w_k^*\}_{k=1}^n = \underset{\{w_k\}}{\operatorname{argmin}} \frac{1}{n} \sum_k \underbrace{\delta(\epsilon_1, \|f \circ g(w_k) - f(x_0)\|_2)}_{\text{boundary constraint}} + \underbrace{\lambda(\epsilon_2 - \bar{d}_n)^+}_{\text{uniformity}}. \quad (2)$$

Examples of generated views are given in Fig. 2.

Scalability of \mathcal{W} -search While effective, online \mathcal{W} -search is computationally expensive because the optimization involving both the contrastive encoder and the generative model needs to be performed for every image in the mini-batch. Therefore, this online view generation using \mathcal{W} -search does not scale as well to large-scale datasets. This scalability problem can be solved by performing view generation *offline*. By leveraging a pretrained f , we can cache the generated views before the actual CL training. In the following experiments, we focus on this offline setting via caching, and provide an empirical study of approximated online version in appendix.

3.2. \mathcal{W} -perturb: perturbation in GM’s latent space

Another solution to the scalability issue of the online \mathcal{W} -search is to generate views via perturbations in the latent space \mathcal{W} of the generator, instead of the latent space of the contrastive encoder \mathcal{Z} . Thereby, we remove the computationally expensive step of finding views through optimization. To avoid on-the-fly optimization, we propose an alternative view generation method, \mathcal{W} -perturb, that creates positive views by directly perturbing in the latent space \mathcal{W} of the pretrained generator g . Under this method, additional positive views for a given anchor x are generated as: $\tilde{x} = \mathcal{W}\text{-perturb}(x) := g(e(x) + w_p)$, where $w_p \sim \mathcal{N}(0, \sigma I)$ and $e(x)$ is the projection of the anchor image in the latent space of g . This is a generalization of the *latent transforms* T_z introduced in [17]. The latent transform is not directly applicable to real image domain since the corresponding latents are unknown. Thus, we project real images in GAN’s latent space via its inverter e .

4. View Assimilation

In the case of SimCLR, prior works [27] have explored replacing expert views entirely with generated views. However, a complete replacement leads to a degradation in performance. Taking the MI maximization perspective of the InfoNCE loss [25, 30, 31, 33] in SimCLR, a possible explanation for the degradation in performance can be attributed to the differences in the MI between the anchor and expert views and the MI between the anchor and generated views. We find that while the original and expert views share roughly the same amount of mutual information as the original and generated views, the shared information seem to be different given that there is a similar gap in information between the expert and the generated views. This finding indicates that the generated views are likely to contain meaningfully *complementary* information to the expert views and hence could lead to additional useful features (for the downstream task). Altogether, these observations motivate us to assimilate generated views into contrastive learning, instead of entirely replacing the expert views, to improve downstream accuracy. To this end, we propose two

methods for *assimilating* generated views into CL training.

Replacement (A1) Our first assimilation method replaces only one of the two expert views with a generated view. On the generated view, we apply a weak amount of *random-resized-crop* and *flipping*.

Multiview (A2) Our second assimilation method simply casts the problem as multiview contrastive learning, where there are more than two positive views. We append the additional positive views, $x_{k(i)}$, and define $\{k(i)\}$ as the batch indices of the appended view(s) generated from an anchor image with the index i . This, however, requires an adjustment in the training loss. To this end, we propose the following multiview loss:

$$L_{\text{multiview}} = L_{\text{CL}} - L_{\text{align}}, \quad \text{where } L_{\text{CL}} = L_{\text{InfoNCE}}$$

$$\text{and } L_{\text{align}} = \sum_{i \in \mathcal{I}} \frac{\alpha}{|k(i)|} \sum_{p \in k(i)} z_i^\top z_p / \tau \quad (3)$$

Our loss function appends an α -weighted sum of dot products of the projected (to the CL embedding space) anchor view i and its respective generated views $k(i)$ to the base CL loss L_{CL} . L_{align} is a general plug-in term that can be used in conjunction with other existing contrastive losses. However, we empirically found it to work best with InfoNCE [31] and adopt it as our L_{CL} unless otherwise specified. Please refer to appendix for results where we use the loss from SimSiam [8]. We also experimented with the multiview loss from SupCon [20] (marked as “A2-full”), but found it to perform more poorly than our proposed loss. From Fig. 3, we can see that our models that use our proposed view generation and assimilation strategies exhibit faster convergence than the baseline SimCLR model.

5. Empirical Study

In this section, we conduct a comprehensive empirical study of view generation and assimilation methods for CL and provide a thorough benchmark of performances. For this purpose, we use the highly optimized SimCLR implementation from [9] for all examined methods and benchmark them on their downstream classification accuracy on four datasets: CIFAR10, CIFAR100 and TinyImageNet. We report the linear probing accuracy as the main evaluation metric. Additional experimental details, hyperparameters, ImageNet experiment and ablations are in appendix.

5.1. View Generation and Assimilation

We start with our main results, ablating all possible combinations of four different view generation methods and two

View 1	View 2	View 3	Loss	CIFAR10	CIFAR100	TinyImageNet	Avg Rank
expert	expert	✗	SimCLR	92.04	70.41	47.48	4.67
expert	\mathcal{W} -search	✗	A1	91.86	71.69	51.08	2.67
expert	\mathcal{W} -perturb	✗	A1	91.09	70.83	50.18	4.67
expert	ViewMaker	✗	A1	82.91	41.87	26.40	8.00
ViewMaker	ViewMaker	✗	SimCLR	83.59	44.04	40.53	7.00
expert	expert	expert	A2	91.46	70.76	47.19	5.33
expert	expert	\mathcal{W} -search	A2	92.90	<u>72.76</u>	<u>51.05</u>	1.67
expert	expert	\mathcal{W} -perturb	A2	<u>92.38</u>	72.95	50.73	<u>2.00</u>
expert	expert	ViewMaker	A2	80.07	36.51	25.30	9.00

Table 1. Linear probe accuracy for the four view generation methods (\mathcal{W} -search, \mathcal{W} -perturb, Viewmaker [27], expert transformation) under A1 and A2 view assimilation methods. We also report the baseline SimCLR (reproduced) and Viewmaker (reproduced) accuracies in rows 1 and 5. The highest Acc@1 for each dataset is highlighted in **boldface**, while the runner-up is underlined.

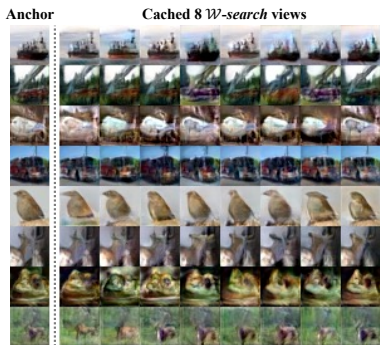


Figure 2. Visual samples of \mathcal{W} -search for CIFAR10 dataset. For each row, we show 8 cached views ($\epsilon_1 = 0.3$).

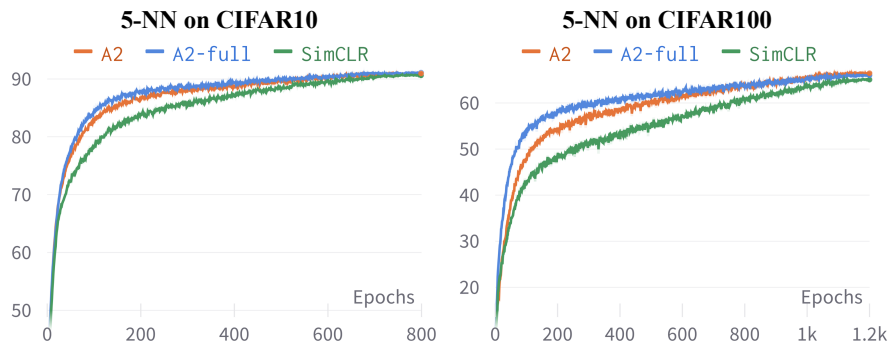


Figure 3. 5-NN accuracy curves during training. Models are trained with \mathcal{W} -search as the view generation and multiview (A2) as the loss. It is clear that the convergence speed of our models is much faster than that of the baseline SimCLR.

view assimilation methods. For the view generation methods, we use our proposed \mathcal{W} -search and \mathcal{W} -perturb methods along with Viewmaker [27] and the expert transformations from SimCLR [4]. For assimilation of generated views, we consider replacement of one expert view (A1) and multiview (A2) as described in Sec. 4. For evaluation, we report the top-1 linear probe accuracy (denoted as Acc@1) and k -Nearest-Neighbor accuracy ($k = 5$, denoted as 5-NN). To obtain linear probe accuracy, we freeze the backbone of f and train a linear layer with SGD for 100 epochs. To determine the value of ϵ 's for \mathcal{W} -search, we first pretrain a SimCLR encoder using expert views and compute the ϵ as the average distance (in \mathcal{W} -space) between anchors and their expertly transformed views. For \mathcal{W} -perturb, we conduct a grid search on the σ .

As shown in Tab. 1, both \mathcal{W} -search and \mathcal{W} -perturb outperform all other view generation methods on CIFAR10, CIFAR100, and TinyImageNet. When we replace one of the views with our proposed view generation strategies, except in CIFAR10, we see consistent improvements and \mathcal{W} -search proves to be a more effective generation method. Especially for TinyImageNet, we see an improvement of $\approx 3.6\%$. When we augment the generated views for mul-

tiview contrastive learning, in contrary to the intuition that more expert views should improve performance, assimilating a third expert view in fact degrades performance in most cases. On the other hand, the views we generate with \mathcal{W} -search and \mathcal{W} -perturb consistently lead to improvements of 0.9%, 2.3%, and 3.6% on CIFAR10, CIFAR100 and TinyImageNet respectively. Overall, our generated views, when replacing one view or being assimilated, almost always leads to an improvement suggests that our framework for view generation allows for generating views that capture some different information from the expert transformations for the downstream task.

6. Conclusion and Limitations

In this work, we presented an empirical study on the view generation and view assimilation techniques in contrastive learning. We showed that when used in conjunction with expert-views, generated views consistently improve downstream classification performance on three different datasets. However, the improvement in the performance is contingent on not only the method of view-generation but more heavily on how the generated view is assimilated, which has been not explored before to our knowledge.

References

- [1] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*, 2018. [8](#)
- [2] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. [1](#)
- [3] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. [1](#), [2](#)
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [4](#)
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning (ICML)*, 2020. [1](#), [2](#)
- [6] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. [1](#), [2](#)
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [1](#)
- [8] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15750–15758, 2021. [1](#), [3](#), [8](#)
- [9] Rumen Dangovski, Li Jing, Charlotte Loh, Seungwook Han, Akash Srivastava, Brian Cheung, Pulkit Agrawal, and Marin Soljačić. Equivariant contrastive learning. *arXiv preprint arXiv:2111.00899*, 2021. [3](#), [7](#)
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [2](#), [7](#)
- [11] Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. *Advances in neural information processing systems*, 32, 2019. [7](#), [10](#)
- [12] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in neural information processing systems (NIPS)*, pages 766–774, 2014. [1](#)
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. [8](#)
- [14] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. [9](#)
- [15] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. [1](#)
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9729–9738, 2020. [1](#)
- [17] Ali Jahanian, Xavier Puig, Yonglong Tian, and Phillip Isola. Generative models as a data source for multiview representation learning. *arXiv preprint arXiv:2106.05258*, 2021. [2](#), [3](#), [7](#), [9](#), [10](#)
- [18] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. [8](#)
- [19] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. [2](#)
- [20] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020. [2](#), [3](#), [7](#)
- [21] Klemen Kotar, Gabriel Ilharco, Ludwig Schmidt, Kiana Ehsani, and Roozbeh Mottaghi. Contrasting contrastive self-supervised representation learning pipelines. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9949–9959, 2021. [1](#)
- [22] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [7](#)
- [23] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. [7](#)
- [24] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6706–6716, 2020. [2](#)
- [25] Ben Poole, Sherjil Ozair, Aaron Van Den Oord, Alex Alemi, and George Tucker. On variational bounds of mutual information. In *International Conference on Machine Learning*, pages 5171–5180. PMLR, 2019. [3](#), [7](#)
- [26] Yuge Shi, N Siddharth, Philip Torr, and Adam R Kosiorek. Adversarial masking for self-supervised learning. In *International Conference on Machine Learning*, pages 20026–20040. PMLR, 2022. [2](#)
- [27] Alex Tamkin, Mike Wu, and Noah Goodman. Viewmaker networks: Learning views for unsupervised representation learning. *arXiv preprint arXiv:2010.07432*, 2020. [1](#), [2](#), [3](#), [4](#), [7](#)

- [28] Ajinkya Tejankar, Soroush Abbasi Koohpayegani, Vipin Pillai, Paolo Favaro, and Hamed Pirsiavash. Isd: Self-supervised learning by iterative similarity distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9609–9618, 2021. [1](#)
- [29] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, 2020. [2](#), [7](#)
- [30] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? *Advances in Neural Information Processing Systems*, 33:6827–6839, 2020. [3](#), [7](#)
- [31] Aaron Van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv e-prints*, pages arXiv–1807, 2018. [3](#), [7](#)
- [32] Guangrun Wang, Keze Wang, Guangcong Wang, Phillip HS Torr, and Liang Lin. Solving inefficiency of self-supervised representation learning. *arXiv preprint arXiv:2104.08760*, 2021. [1](#)
- [33] Mike Wu, Chengxu Zhuang, Milan Mosse, Daniel L. K. Yamins, and Noah D. Goodman. On mutual information in contrastive learning for visual representations. *ArXiv*, abs/2005.13149, 2020. [3](#), [7](#)
- [34] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3733–3742, 2018. [1](#)
- [35] Mang Ye, Xu Zhang, PongChi Yuen, and Shih-Fu Chang. Unsupervised embedding learning via invariant and spreading instance feature. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6203–6212, 2019. [2](#)
- [36] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning (ICML)*, 2021. [1](#)
- [37] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *European conference on computer vision*, pages 592–608. Springer, 2020. [8](#)
- [38] Roland S Zimmermann, Yash Sharma, Steffen Schneider, Matthias Bethge, and Wieland Brendel. Contrastive learning inverts the data generating process. In *International Conference on Machine Learning*, pages 12979–12990. PMLR, 2021. [1](#), [2](#)

Appendix

A. Experimental Setup and Details

Datasets. Experiments are conducted on four datasets:

- CIFAR10 [22] has 10 classes, 50,000 images for training and 10,000 for testing.
- CIFAR100 [22] has 100 classes, 50,000 images for training and 10,000 for testing.
- TinyImageNet is introduced in [23]. The dataset contains 200 classes, 100,000 images for training and 10,000 for testing. Images are resized to 64×64 .
- ImageNet [10] contains approximately 1.3 million images. We following the setting in [17] and generate 1.3 million “fake” images using a pretrained BigBiGAN [11] at resolution 128×128 .

Implementation. We implement our methods based on the E-SSL [9] codebase¹ (for CIFAR10 experiments) and the SupCon [20] codebase² (for CIFAR100 and TinyImageNet). Hyperparameters for each dataset are listed in Tab. 2³.

Viewmaker. For Viewmaker [27], we reproduced the reported accuracy on CIFAR10. For CIFAR100 and TinyImageNet, since the original authors did not evaluate their model against these datasets, we tried our best to optimize the hyperparameters, such as optimizer, learning rate, temperature, architecture of encoder (ResNet18 small, ResNet18, ResNet50), and projection head. We describe the best set of hyperparameters in Tab. 3.

Evaluation. As for evaluation metrics, we adopt the conventions in the respective codebases. For CIFAR10, we run linear probe evaluations (for 100 epochs) with 5 random seeds and report the mean and standard deviation of accuracies. For CIFAR100 and TinyImageNet, we run linear probe evaluations for 100 epochs and report the best accuracy. For both settings, we load and freeze the last checkpoint of the backbone network.

Computational cost for view generation. The computation time depends on the hyperparameters (the number of optimization steps for \mathcal{W} -search), e.g., caching 8 views per sample for CIFAR10 takes 12.19 A100 GPU hours.

Computational cost for pretraining. Each experiment is run on 4 NVIDIA V100 GPUs. The pretraining time of SimCLR baseline for CIFAR10, CIFAR100, and TinyImageNet are 11.5, 13.6, and 29.7 hours, respectively.

¹<https://github.com/rdangovs/essl/tree/main/cifar10>

²<https://github.com/HobbitLong/SupContrast>

³For CIFAR10, we found that batch size of 128 gives similar or slightly better results than the default 512.

	CIFAR10	CIFAR100	TinyImageNet
Optimizer	SGD	SGD	SGD
Learning Rate	0.015	0.5	0.5
Weight Decay	5e-5	1e-4	1e-4
Momentum	0.9	0.9	0.9
Cosine Decay	✓	✓	✓
Batch Size	128	512	512
SimCLR Loss	InfoNCE	SimCLR	SimCLR
Temperature	0.5	0.5	0.5
Epochs	800	1200	1000
Backbone	ResNet18	ResNet50	ResNet18
Embedding Dim	512	2048	512
Projection Dim	2048	128	128

Table 2. Hyperparameters for experiments.

	CIFAR10	CIFAR100	TinyImageNet
Optimizer	SGD (for encoder), Adam (for Viewmaker module)		
Learning Rate	0.015	0.06	0.06
Weight Decay	1e-4	1e-4	1e-4
Momentum	0.9	0.9	0.9
Cosine Decay	✗	✗	✗
Batch Size	128	512	128
SimCLR Loss	ViewMaker	ViewMaker	ViewMaker
Temperature	0.07	0.1	0.5
α for A2 loss	0.14	0.1	0.5
Epochs	200	800	800
Backbone	ResNet18	ResNet50	ResNet18
Noise Dim	100	100	100
Embedding Dim	512	2048	512
Projection Dim	128	128	128

Table 3. Hyperparameters for our reproduced ViewMaker [27] experiments.

B. Mutual Information Analysis

Taking the mutual information (MI) maximization perspective of the InfoNCE [25, 30, 31, 33] loss in SimCLR, [29, 30] introduced an information theoretic definition of what makes for a “good” positive view in contrastive learning. While it is difficult to accurately estimate MI in the high-dimensional \mathcal{X} space, in Table 4, we provide rough estimates for initial analysis. We find that the original and generated views share similar or lower amount of mutual information than the original and expert views. However, generated (\mathcal{W} -search or \mathcal{W} -perturb) and expert views share even lower mutual information than original and expert views, which indicates that VG1 views are likely to contain meaningfully complementary information to the expert views and hence lead to additional useful features (for the downstream task). These observations motivate us to assimilate generated views into contrastive learning, instead of entirely replacing the expert views, to improve downstream accuracy.

View Pairs	CIFAR10	CIFAR100
Original, Expert	4.14	5.41
Original, \mathcal{W} -search	4.13	4.40
\mathcal{W} -search, Expert	3.78	4.35
Original, \mathcal{W} -perturb	3.79	3.91
\mathcal{W} -perturb, Expert	3.66	3.86

Table 4. Estimated mutual information between the anchor image (original), positive expert view (expert) and generated views (\mathcal{W} -search and \mathcal{W} -perturb) using the MINE [1] estimator.

C. Details of StyleGAN and In-Domain GAN Inversion

StyleGAN We use StyleGAN2⁴ for our experiments on CIFAR and TinyImageNet. The StyleGAN generator consists of two key components: (1) a mapping function $g_1 : \mathcal{S} \mapsto \mathcal{W}$ that maps the Gaussian-distributed latent code $s \in \mathcal{S}$ into a collection of style codes $w \in \mathcal{W}$, and (2) a generator $g : \mathcal{W} \mapsto \mathcal{X}$ that decodes $w \in \mathcal{W}$ to an image. Here, w is a concatenation of w_1, \dots, w_k , where each w_i corresponds to the style code from the i^{th} convolutional block of g .

In-Domain GAN Inversion Let $e : \mathcal{X} \mapsto \mathcal{W}$ denote an inverter neural network. Let $d : \mathcal{X} \mapsto \mathbb{R}$ denote the discriminator network. In-domain GAN inversion [37] aims to learn a mapping from images to latent space. The encoder is trained to reconstruct real images (thus are ‘‘in-domain’’) and guided by image-level loss terms, *i.e.* pixel MSE, VGG perceptual loss, and discriminator loss:

$$L_{\text{idinv}}(e, d, g) = \mathbb{E}_{x \sim P_X} [\|x - g \circ e(x)\|_2 + \lambda_{\text{vgg}} \|h(x) - h \circ g \circ e(x)\|_2 - \lambda_{\text{adv}} \mathcal{A}(-\tilde{d} \circ g \circ e(x))], \quad (4)$$

where h is perception network and here we keep the same as in-domain inversion as VGG network, a is the *activation function* and \tilde{d} is the *logit* or discriminator’s output before activation. Note that choosing $a(t) = \text{softplus}(t) = \log(1 + \exp(t))$ recovers the original GAN formulation [13, 18], and the resulting objective minimizes the Jensen-Shannon divergence between real and generated data distributions. After encoder training, we optimize the associated latent variable w for each image x with the same loss function using $w = e(x)$ as a warm start. Note that in the main text we reload the notation $e(\cdot)$ as the final results after w -optimization, which are precomputed and cached.

D. Details of Training Loss

We provide an illustrative visualization of our A2 losses in Fig. 4. The detailed formulation of loss functions are as follows,

$$L_{\text{infonce}} = - \sum_{i \in \mathcal{I}_1} \log \frac{\exp(z_i \cdot z_{j(i)}/\tau)}{\sum_{a \in \mathcal{I}_2} \exp(z_i \cdot z_a/\tau)} - \sum_{i \in \mathcal{I}_2} \log \frac{\exp(z_i \cdot z_{j(i)}/\tau)}{\sum_{a \in \mathcal{I}_1} \exp(z_i \cdot z_a/\tau)} \quad (5)$$

$$L_{\text{A2-full}} = - \sum_{i \in \mathcal{I}} \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(z_i \cdot z_p/\tau)}{\sum_{a \in \mathcal{A}(i)} \exp(z_i \cdot z_a/\tau)} \quad (6)$$

where $P(i) = \{j(i)\} \cup \{k(i)\}$

$$L_{\text{A2-simclr}} = L_{\text{simclr}} - \sum_{i \in \mathcal{I}} \frac{\alpha}{|k(i)|} \sum_{p \in k(i)} z_i \cdot z_p/\tau = - \sum_{i \in \mathcal{I}} \log \frac{\exp(z_i \cdot z_{j(i)}/\tau + \frac{\alpha}{|k(i)|} \sum_{p \in k(i)} (z_i \cdot z_p)/\tau)}{\sum_{a \in \mathcal{A}(i)} \exp(z_i \cdot z_a/\tau)} \quad (7)$$

In L_{infonce} , \mathcal{I}_1 and \mathcal{I}_2 are the set of indices of two positive views. An ablation of A2 losses is provided in Tab. 5. We observe that A2-InfoNCE performs the best for both datasets. We used A2-InfoNCE as our A2 loss if not specified.

Loss	CIFAR10		CIFAR100	
	Acc@1	5-NN	Acc@1	5-NN
A2-full	92.57	91.57	71.82	66.06
A2-SimCLR ($\alpha = 0.5$)	92.66	91.05	72.27	65.85
A2-InfoNCE ($\alpha = 0.5$)	92.90	90.95	72.76	66.46
A2-InfoNCE ($\alpha = 1$)	92.54	90.80	72.61	66.96

Table 5. Ablation on the loss functions with VG1 views on the CIFAR10 and CIFAR100 datasets. A2-SimCLR is $L_{\text{simclr}} - L_{\text{align}}$ and A2-InfoNCE is $L_{\text{infonce}} - L_{\text{align}}$.

E. SimSiam Experiments

We conducted experiments of SimSiam [8] with generated views. The 5-NN accuracy curves are reported in Fig. 5. The linear probe accuracy is reported in Tab. 6. The A2 loss needs to be adjusted accordingly to avoid training collapse,

$$L_{\text{A2-simsiam}} = L_{\text{simsiam}} + \sum_{i \in \mathcal{I}} \frac{\alpha}{|k(i)|} \sum_{p \in k(i)} D(\text{predictor}(z_p), \text{stopgrad}(z_i)) \quad (8)$$

where D is the cosine similarity and L_{simsiam} is the SimSiam loss function.

⁴<https://github.com/rosinality/stylegan2-pytorch>

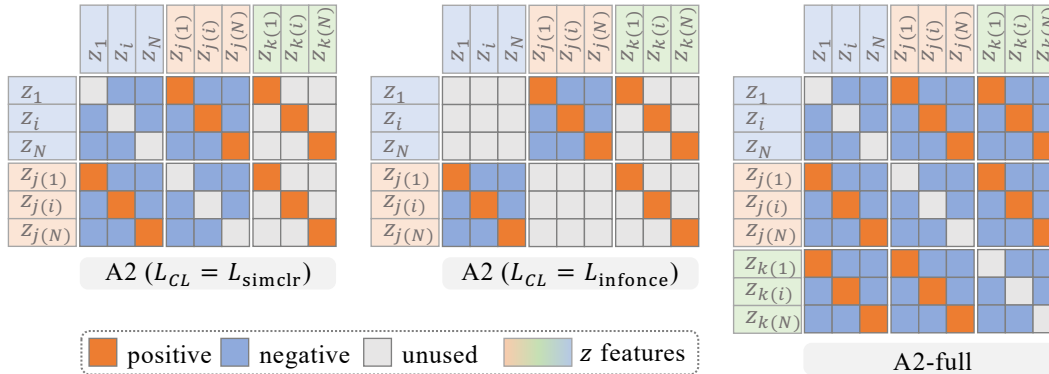


Figure 4. Visual illustration of the three contrastive loss functions: A2-SimCLR, A2-InfoNCE, and A2-full for a mini-batch size of N images (visualizing $N = 3$). Each small square represents the inner-product between corresponding features. The contrastive loss for each instance i is defined on the corresponding row of inner-products.

	Acc@1	5-NN
SimSiam baseline	90.17	89.27
\mathcal{W} -search-SimCLR + A2-SimSiam	90.96	89.81
\mathcal{W} -search-SimSiam + A2-SimSiam	90.79	89.82
\mathcal{W} -perturb + A2-SimSiam	90.28	89.99
SimCLR baseline	92.04	90.65
\mathcal{W} -search-SimSiam + A2-InfoNCE	92.44	90.88

Table 6. Linear-probing and 5-NN accuracies of SimSiam experiments on CIFAR10.

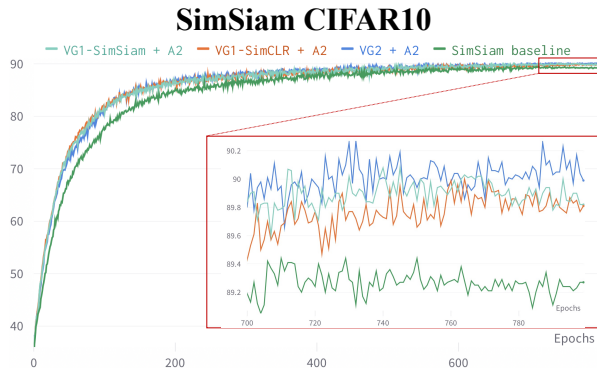


Figure 5. 5-NN accuracy curves of SimSiam experiments on CIFAR10. In legend, VG1 refers to \mathcal{W} -search and VG2 refers to \mathcal{W} -perturb.

F. Online \mathcal{W} -search

In the online \mathcal{Z} -search setting, the optimization is performed involving the current SimCLR encoder during training. We tried to perform 1-step optimization with fast sign gradient [14], but observed the results are worse than the SimCLR baseline. The 5-NN accuracies are reported in Fig. 6.

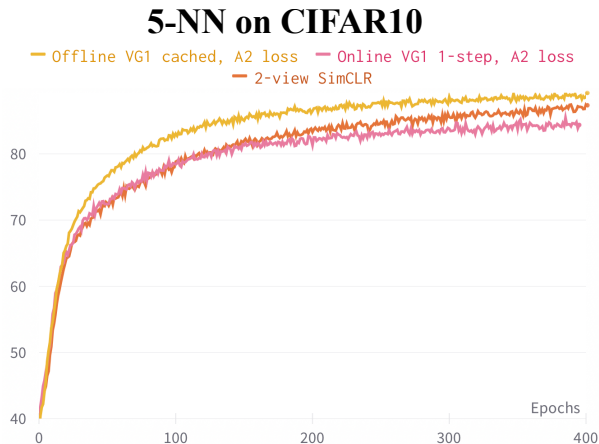


Figure 6. Online \mathcal{W} -search. In legend, VG1 refers to \mathcal{W} -search.

G. Ablation on Hyperparameters

In this section we provide ablations on hyperparameters ϵ_1 , ϵ_2 , and λ introduced in \mathcal{Z} -search. In addition, we perform grid search on σ introduced in \mathcal{Z} -perturb.

Ablation on ϵ_1 and ϵ_2 . We conduct ablation studies of ϵ_1 , ϵ_2 , and λ on CIFAR10. In Tab. 9, we set $\epsilon_2 = \epsilon_1 + 0.2$ except for ϵ_1 of values 0.1 and 0.2. We empirically find that it is difficult to reach a large pairwise distance ϵ_2 when ϵ_1 is small, and a large ϵ_2 leads to more optimization steps. By design, a large ϵ_2 encourages generating diverse samples. A rule of thumb is to set $\epsilon_2 \geq \epsilon_1$.

Ablation on λ . In Tab. 10, we fix $\epsilon_2 = 0.15$ and vary ϵ_1 . In Tab. 11, we fix $\epsilon_1 = 0.3$ and $\epsilon_2 = 0.5$ and vary λ .

Ablation on σ . We perform grid search on σ for VG2 and report results in Tab. 7 and Tab. 8. We find that for both CIFAR10 and CIFAR100, $\sigma = 0.2$ leads to the best results, which is consistent with the empirical findings in [17].

View	Acc@1
$e(x)$	91.697 \pm 0.038
$e(x) + w_{\text{Gauss}} \sim N(0, 0.1)$	92.297 \pm 0.024
$e(x) + w_{\text{Gauss}} \sim N(0, 0.2)$	92.383 \pm 0.033
$e(x) + w_{\text{Gauss}} \sim N(0, 0.4)$	91.852 \pm 0.030
$e(x) + w_{\text{Gauss}} \sim N(0, 1.0)$	87.893 \pm 0.026

Table 7. Ablation on σ of value 0, 0.1, 0.2, 0.4 and 1.0. Experiments are conducted on CIFAR10 with A2 loss.

View	Acc@1	5-NN
$e(x) + w_{\text{Gauss}} \sim N(0, 0.1)$	71.84	66.33
$e(x) + w_{\text{Gauss}} \sim N(0, 0.2)$	72.95	66.60
$e(x) + w_{\text{Gauss}} \sim N(0, 0.4)$	71.69	65.33
$e(x) + w_{\text{Gauss}} \sim N(0, 1.0)$	68.05	61.06

Table 8. Ablation on σ of value 0.1, 0.2, 0.4 and 1.0. Experiments are conducted on CIFAR100 with A2 loss.

ϵ_1	ϵ_2	Acc@1
0.1	0.15	92.584 \pm 0.023
0.2	0.35	92.615 \pm 0.048
0.3	0.50	92.898 \pm 0.045
0.5	0.70	92.451 \pm 0.053
0.7	0.90	91.760 \pm 0.032
0.9	1.10	91.561 \pm 0.051

Table 9. Ablation on ϵ_1 and ϵ_2 , $\lambda = 0.01$. Experiments are conducted on CIFAR10.

ϵ_1	ϵ_2	Acc@1
0.1	0.15	92.584 \pm 0.023
0.2	0.15	92.385 \pm 0.043
0.3	0.15	92.643 \pm 0.043
0.5	0.15	92.455 \pm 0.032
0.7	0.15	92.174 \pm 0.029

Table 10. Ablation on ϵ_1 , $\epsilon_2 = 0.15$ and $\lambda = 0.01$. Experiments are conducted on CIFAR10 with the A2 loss.

ϵ_1	λ	Acc@1
0.3	0	92.555 \pm 0.028
0.3	0.005	92.756 \pm 0.042
0.3	0.01	92.898 \pm 0.045
0.3	0.02	92.854 \pm 0.027

Table 11. Ablation studies on λ . For all entries we fix $\epsilon_1 = 0.3$ and $\epsilon_2 = 0.5$. Experiments are conducted on CIFAR10 with the A2 loss.

View 1	View 2	View 3	Loss	Acc@1
expert	expert	X	SimCLR	49.93
expert	\mathcal{W} -perturb	X	A1	48.81
expert	expert	\mathcal{W} -perturb	A2	51.42

Table 12. ImageNet experiments. Linear probe accuracies are reported.

H. ImageNet Experiments

We also evaluate against the standard large-scale dataset ImageNet. Following the more difficult setting in [17], we generate 1.3 million “fake” images using BigBiGAN [11] (approximately the number of images in ImageNet) from the GAN and only train on this generated dataset. However, we follow the standard protocol of reporting the linear probe accuracy on the real ImageNet dataset. As evident in Tab. 12, adding our generated view as an additional positive view improves performance by $\approx 1.5\%$ and proves to capture some meaningful information that expert transformations do not.