

---

# Stochastic Contrastive Learning

---

Jason Ramapuram\*, Dan Busbridge\*, Xavier Suau, Russ Webb  
Apple  
{jramapuram, dbusbridge, xsuaucuadros, rwebb}@apple.com

## Abstract

While state-of-the-art contrastive Self-Supervised Learning (SSL) models produce results competitive with their supervised counterparts, they lack the ability to infer latent variables. In contrast, prescribed latent variable (LV) models enable attributing uncertainty, inducing task specific compression, and in general allow for more interpretable representations. In this work, we introduce LV approximations to large scale contrastive SSL models. We demonstrate that this addition improves downstream performance (resulting in 96.42% and 77.49% test top-1 fine-tuned performance on CIFAR10 and ImageNet respectively with a ResNet50) as well as producing highly compressed representations (588× reduction) that are useful for interpretability, classification and regression downstream tasks.

## 1 Introduction

Learning meaningful representations without human domain knowledge has been a long-standing goal of machine learning. Recent work in large scale SSL (Chen et al., 2020a,b; Alayrac et al., 2020; Grill et al., 2020; Caron et al., 2020; Zbontar et al., 2021; Caron et al., 2021) has advanced this pursuit and narrowed the gap against fully supervised models, all the while relaxing the use of potentially biased human labels. And yet, the SSL methods and toolkits lack a method to add interpretable, prescribed distributions into the representation learning process. In this work, we address this shortcoming through the introduction of Bernoulli and Isotropic-Gaussian latent variables into the SimCLR (Chen et al., 2020a) contrastive learning framework.

The use of Bernoulli latent variables enables extracting meaningful discrete representations of image data, providing a natural means of data dependent compression<sup>2</sup> that is useful for downstream tasks such as classification and regression. Interestingly, we find that the use of discrete latent variables improves downstream performance when fully finetuning the representation learning backbone, outperforming SimCLR (Chen et al., 2020a) on CIFAR10 and ImageNet1000 (Deng et al., 2009).

## 2 Background

In this work we focus on large scale contrastive learning, where we optimize the InfoNCE objective (van den Oord et al., 2018; Chen et al., 2020a). InfoNCE generalizes Noise Contrastive Estimation (NCE) by using variates from the empirical data distribution,  $\{\mathbf{x}_i, \mathbf{x}_j\} \sim p(\mathbf{x})$ , mapping them through networks,  $g_\theta$  and  $f_\theta$ , to a representation  $\mathbf{v} = (g_\theta \circ f_\theta)(\mathbf{x})$ .  $f_\theta$  is typically referred to as the backbone and  $g_\theta$  as the InfoNCE head. While NCE samples negative variates from a naive prior,  $\mathbf{v} \sim p(\mathbf{v})$ , InfoNCE uses true variates in a multi-sample un-normalized bound: (Poole et al., 2019)

$$\mathcal{L}_{\text{InfoNCE}}^{(i,j)} = -\log \frac{\exp(\text{sim}(\mathbf{v}_i, \mathbf{v}_j)/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(\text{sim}(\mathbf{v}_i, \mathbf{v}_k)/\tau)}. \quad (1)$$

---

\*Equal contribution.

<sup>2</sup>The map  $\mathbb{R}^{224 \times 224 \times 3} \mapsto \mathbb{Z}_2^{2048}$  implies a  $\frac{1}{588} \times$  compression given 8 bits/input pixel and 1 bit/binary output.

The similarity operator (sim) from Equation 1 typically is modeled with a cosine-similarity on the representation feature vectors  $\{v_i, v_j\}$ , with a controllable temperature hyper-parameter  $\tau$ .

### 3 Stochastic Contrastive Learning (StochCon)




Figure 1: StochCon without multi-crop (Caron et al., 2020) using plate notation. Following Chen et al. (2020a), two data augmentation operators  $t, t'$  are sampled from the augmentation family  $\mathbb{T}$ , producing views  $\tilde{x}, \tilde{x}'$ . A backbone network  $f$  is applied to each example, producing  $\mathbf{h}, \mathbf{h}' \in \mathbb{R}^{D_{\text{Backbone}}}$  (we include spatial pooling in our definition of  $f$ ). On the bottom branch, the  $\mathbf{h}'$  specifies the natural parameters of a distribution  $q(\mathbf{z}; \pi(\mathbf{h}'))$ , via an encoder  $\pi : \mathbb{R}^{D_{\text{Backbone}}} \rightarrow \mathbb{R}^{D_{\text{Latent}}}$ . Variates  $\mathbf{z}' \sim q$  are decoded by  $\rho : \mathbb{R}^{D_{\text{Latent}}} \rightarrow \mathbb{R}^{D_{\text{Backbone}}}$  and compared using the standard  $\mathcal{L}_{\text{InfoNCE}}$  (Equation 1) using the head  $g$ . This work’s contribution is the stochastic mapping  $\phi' \mapsto \mathbf{z}' \mapsto \mathbf{h}''$ . Gray circles denote random variables. Bottom branch negative variates are reparameterized in the same manner.

---

#### Algorithm 1 StochCon

---

**Require:** Data:  $x \sim p(x), t \sim \mathbb{T}(x)$

**Require:** Models:  $f_\theta$  : backbone,  $g_\theta$  : head,  $\{\pi_\theta, \rho_\theta\}$  : projectors

**while** not converged **do**

$\{\hat{x}, \hat{x}'\} = \{t \circ x, t' \circ x\}$

▸ Augment input with  $\{t, t'\}$

$\{\mathbf{h}, \mathbf{h}'\} = \{f_\theta(\hat{x}), f_\theta(\hat{x}')\}$

▸ Produce representations

$\phi' = \pi_\theta(\mathbf{h}')$

▸ (optional) Bottleneck projection

$\mathbf{z}' \sim q_\theta(\mathbf{z}|\mathbf{x})$

▸ Pathwise differentiable (Mohamed et al., 2020) latent variable.

$\mathbf{h}'' = \rho_\theta(\mathbf{z}')$

▸ (optional) Bottleneck upsampler

$\{v, v'\} = \{g_\theta(\mathbf{h}), g_\theta(\mathbf{h}'')\}$

▸ InfoNCE projection

$\min_\theta \mathcal{L}_{\text{InfoNCE}}(v, v')$

**end while**

---

We describe our model in Algorithm 1 and Figure 1. We modify SimCLR (Chen et al., 2020a) by forcing bottom branch variates through a pathwise differentiable (Mohamed et al., 2020) distribution,  $\mathbf{z}' \sim q_\theta(\mathbf{z}|\mathbf{x})$ . Importantly,  $\mathbf{z}'$  can be optionally projected to a lower dimensional space,  $|\mathbf{z}'| \ll |\mathbf{h}|$ , through linear projection layers,  $\{\pi_\theta, \rho_\theta\}$ . Upon ablation, we observe minimal performance degradation when projecting one branch of the SimCLR model through a low dimensional distribution, with the advantage of having more interpretable features (Section 4.1).

In this work, we explore the isotropic-Gaussian (Kingma & Welling, 2014) and Gumbel-Bernoulli (Jang et al., 2017; Maddison et al., 2017) distributions. We apply the differentiable distribution on the output of the backbone model  $f_\theta$ , given an optional bottleneck projection  $\pi_\theta$ .

## 4 Experiments

**Training details** Following Chen et al. (2020a), all models train with a batch size of 4096, the LARS optimizer (Huo et al., 2021) with linear warmup (Goyal et al., 2017) and a single cycle cosine annealed learning rate schedule (Goyal et al., 2017; Smith & Topin, 2017). We use DINO augmentations (Caron et al., 2021) (2-global views + 8-local views (Caron et al., 2020)) for all SimCLR

variants. For the Gumbel-Bernoulli distribution, the temperature is annealed from 1.0  $\rightarrow$  0.1 using a single cycle cosine schedule during training. Finetuning procedure is described in Appendix A.3.

Model performance for linear-probes on a non-updated (*Frozen*) and fine-tuned (*Fine-Tuned*) backbone is given in Table 1. We observe that StochCon (*Fine-Tuned*) outperforms an equally tuned SimCLR model, as well as a supervised model with the same ResNet50 and ResNet200 (He et al., 2016) architectures, while the *Frozen* probe is competitive. We validate that this performance difference does not arise purely from the Gumbel-Bernoulli through ablations presented in Appendix A.2.

Table 1: Summary test top-1% for CIFAR10 and ImageNet1000.

Model	CIFAR10-ResNet50		ImageNet-ResNet50		ImageNet-ResNet200	
	Fine-Tuned	Frozen	Fine-Tuned	Frozen	Fine-Tuned	Frozen
StochCon Bern	<b>96.42</b>	91.96	<b>77.49</b>	67.00	<b>80.24</b>	64.25
StochCon Iso-Gauss	96.08	<b>92.40</b>	–	–	–	–
Supervised	95.00	–	76.13	–	78.34	–
SimCLR	94.35	91.67	76.37	<b>71.34</b>	79.82	<b>73.52</b>

#### 4.1 Ablations

To evaluate the benefits of our StochCon framework, we propose a series of ablations. In Figure 2-*Left*, we train StochCon Bernoulli models with varying bottleneck  $\phi'$  dimensions and present the top-1 *Frozen* performance of each model. Results show that StochCon is robust to variadic sized latents. We believe this robustness is due to the model learning to compare a full  $\mathbf{h} \in \mathbb{R}^{2048}$  dimensional vector to an upsampled low dimensional latent,  $\mathbf{z}' \in \mathbb{R}^D, D \ll 2048$ .

In Figure 2-*Right*, we analyze the mean  $F_1$  performance for a multi-class Random Forest evaluated by varying the number of feature units. Surprisingly, we find that to accurately classify CIFAR10, the StochCon model with a 64 dimensional latent Gumbel-Bernoulli only requires 11 binary feature units. We also observe that performance decreases for the Isotropic-Gaussian as we increase the latent dimensionality, holding the number of Random Forest units constant. Note that this does not happen in the Gumbel-Bernoulli case, as this variable does form distributed representations in the way an Isotropic-Gaussian does.

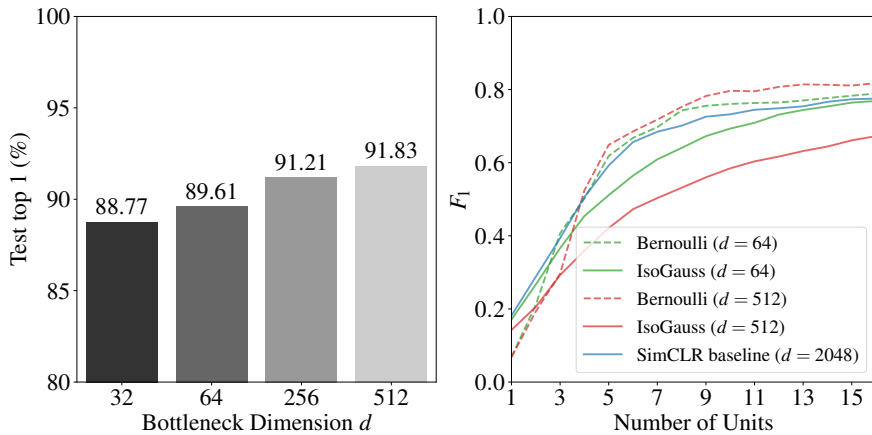


Figure 2: **Left:** Ablation over bottleneck latent dimensions for Bernoulli-StochCon. **Right:** Mean  $F_1$  performance across all classes on CIFAR10 for a multi-class Random Forest, varying the number of feature units. Units are identified on the training set using Random Forest feature importance under 5-fold stratified sampling. Performance is the mean over the held-out test sets.

#### 4.2 Isotropic-Gaussian StochCon and variance collapse

Since StochCon does not constrain the latent variable distribution, we observed that in the case of a learned variance where  $\mathbf{z}' \sim \mathcal{N}(\boldsymbol{\mu}(\boldsymbol{\pi}(\mathbf{h}')), \boldsymbol{\sigma}^2(\boldsymbol{\pi}(\mathbf{h}')))$ , the learned variance,  $\boldsymbol{\sigma}^2(\boldsymbol{\pi}(\mathbf{h}'))$  would triv-

ially collapse to 0. To work around this and provide meaningful uncertainties, we force the network to learn to estimate variances of the opposing set of views, so that  $z' \sim \mathcal{N}(\mu(\pi(\mathbf{h}')), \sigma^2(\pi(\mathbf{h}')))$ . We validate this below in Figure 3 and find that as the bottleneck dimension reduces, the model learns to rely more on the available stochasticity.




Figure 3: CIFAR10 test top-1 (**Left**) and test aggregate variance,  $\text{Var}_x[q_\theta(z|\mathbf{x})]$ , (**Right**) for varying isotropic-Gaussian bottleneck dimensions.

### 4.3 Countable metrics for Bernoulli-StochCon

Since StochCon works with discrete representation vectors, it enables analysis through countable metrics. We present the average count of representation bits for the  $\mathbb{R}^{2048}$  dimensional StochCon Bernoulli model in Figure 4. We ablate four different variants:  $\{hard\ bottom, hard\ top, soft\ bottom, soft\ top\}$ . The difference between these variants is where the distribution is applied: the *top*-\* models apply the reparameterization on the global image views (Caron et al., 2020), while the *bottom*-\* models apply them on the local views (Caron et al., 2020). The *hard*-\* models use a differentiable mechanism to always feed-forward discrete variates (see Appendix Section A.1), while the *soft*-\* models use the standard variates extracted from the Gumbel-Bernoulli distribution.




Figure 4: CIFAR10 average bit counts, aggregated across training (**Left**) and test (**Right**) datasets.

At the beginning of training we observe that the average number of active bits is approximately half of the available  $\mathbb{R}^{2048}$ , but as training progresses this quantity decreases. Note that this does not imply that the model does not use the zero valued bits, but rather provides an alternative method to analyze model performance. For reference in Figure 4-*Right* we also include the test active bits for the 512 dimensional bottleneck model (*soft-top 512*). We observe that when the model is capacity restricted it uses all available Bernoulli latents (50% of its representation for zeros, 50% for ones).

## 5 Conclusion

In this work, we present a novel formulation that enables the use of latent variables in large scale contrastive self-supervised models. We demonstrate that in addition to improving downstream performance these models reveal that in practice competitive discriminative performance on CIFAR10 can be achieved with as few as 11 bits (Figure 2). Future work will explore further latent variable models such as the pathwise Beta distribution (Figurnov et al., 2018) and non-parametric latents such as normalizing flows (Rezende & Mohamed, 2015).

## Acknowledgments

The authors would like to thank the following people for their help throughout the process of writing this paper, in alphabetical order: Barry-John Theobald, Katherine Metcalf, Luca Zappella and Miguel Sarabia del Castillo. Additionally, we thank Andrea Klein, Cindy Liu, Guihao Liang, Guillaume Seguin, Li Li, Okan Akalin, and the wider Apple infrastructure team for assistance with developing scalable, fault tolerant code.

## References

- Jean-Baptiste Alayrac, Adrià Recasens, Rosalia Schneider, Relja Arandjelovic, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/0060ef47b12160b9198302ebdb144dcf-Abstract.html>.
- Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/70feb62b69f16e0238f741fab228fec2-Abstract.html>.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *CoRR*, abs/2104.14294, 2021. URL <https://arxiv.org/abs/2104.14294>.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1597–1607. PMLR, 2020a. URL <http://proceedings.mlr.press/v119/chen20j.html>.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E. Hinton. Big self-supervised models are strong semi-supervised learners. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020b. URL <https://proceedings.neurips.cc/paper/2020/hash/fcbc95ccdd551da181207c0c1400c655-Abstract.html>.
- Ekin D. Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. Randaugment: Practical automated data augmentation with a reduced search space. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2020, Seattle, WA, USA, June 14-19, 2020*, pp. 3008–3017. Computer Vision Foundation / IEEE, 2020. doi: 10.1109/CVPRW50498.2020.00359. URL [https://openaccess.thecvf.com/content\\_CVPRW\\_2020/html/w40/Cubuk\\_Randaugment\\_Practical\\_Automated\\_Data\\_Augmentation\\_With\\_a\\_Reduced\\_Search\\_Space\\_CVPRW\\_2020\\_paper.html](https://openaccess.thecvf.com/content_CVPRW_2020/html/w40/Cubuk_Randaugment_Practical_Automated_Data_Augmentation_With_a_Reduced_Search_Space_CVPRW_2020_paper.html).

- Ekin Dogus Cubuk, Barret Zoph, Dandelion Mané, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation policies from data. *CoRR*, abs/1805.09501, 2018. URL <http://arxiv.org/abs/1805.09501>.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pp. 248–255. IEEE Computer Society, 2009. doi: 10.1109/CVPR.2009.5206848. URL <https://doi.org/10.1109/CVPR.2009.5206848>.
- Mikhail Figurnov, Shakir Mohamed, and Andriy Mnih. Implicit reparameterization gradients. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 439–450, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/92c8c96e4c37100777c7190b76d28233-Abstract.html>.
- Priya Goyal, Piotr Dollár, Ross B. Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch SGD: training imagenet in 1 hour. *CoRR*, abs/1706.02677, 2017. URL <http://arxiv.org/abs/1706.02677>.
- Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent - A new approach to self-supervised learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/f3ada80d5c4ee70142b17b8192b2958e-Abstract.html>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- Zhouyuan Huo, Bin Gu, and Heng Huang. Large batch optimization for deep learning using new complete layer-wise adaptive rate scaling. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pp. 7883–7890. AAAI Press, 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16962>.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=rkE3y85ee>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *5th International Conference on Learning Representations*,

*ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=S1jE5L5g1>.

Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. Monte carlo gradient estimation in machine learning. *J. Mach. Learn. Res.*, 21:132:1–132:62, 2020. URL <http://jmlr.org/papers/v21/19-346.html>.

Ben Poole, Sherjil Ozair, Aaron van den Oord, Alex Alemi, and George Tucker. On variational bounds of mutual information. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5171–5180. PMLR, 2019. URL <http://proceedings.mlr.press/v97/poole19a.html>.

Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In Francis R. Bach and David M. Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 1530–1538. JMLR.org, 2015. URL <http://proceedings.mlr.press/v37/rezende15.html>.

Leslie N. Smith and Nicholay Topin. Super-convergence: Very fast training of residual networks using large learning rates. *CoRR*, abs/1708.07120, 2017. URL <http://arxiv.org/abs/1708.07120>.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *CoRR*, abs/1807.03748, 2018. URL <http://arxiv.org/abs/1807.03748>.

Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. *CoRR*, abs/2103.03230, 2021. URL <https://arxiv.org/abs/2103.03230>.

## A Appendix

### A.1 Discrete Gumbel-Bernoulli variates

The Gumbel-Bernoulli distribution in its naive form returns non discretized variates when the temperature,  $\tau$ , is high. However, a well known trick to extract proper discrete variates is summarized in the pytorch code below.

```
1 import torch
2
3 def compute_hard(relaxed: torch.Tensor) -> torch.Tensor:
4     """Produce a hard version of relaxed as a differentiable tensor.
5
6     :param relaxed: the relaxed estimate
7     :returns: hard bernoulli with 0s and 1s
8     :rtype: torch.Tensor
9
10    """
11    hard = relaxed.clone()
12    hard[relaxed < 0.5] = 0.0
13    hard[relaxed >= 0.5] = 1.0
14    hard_diff = hard - relaxed # sub the relaxed tensor backprop path
15    return hard_diff.detach() + relaxed # add back to keep bp path
```

Listing 1: Hard Gumbel-Bernoulli variates.

### A.2 SimCLR Finetuning and Supervised Bernoulli

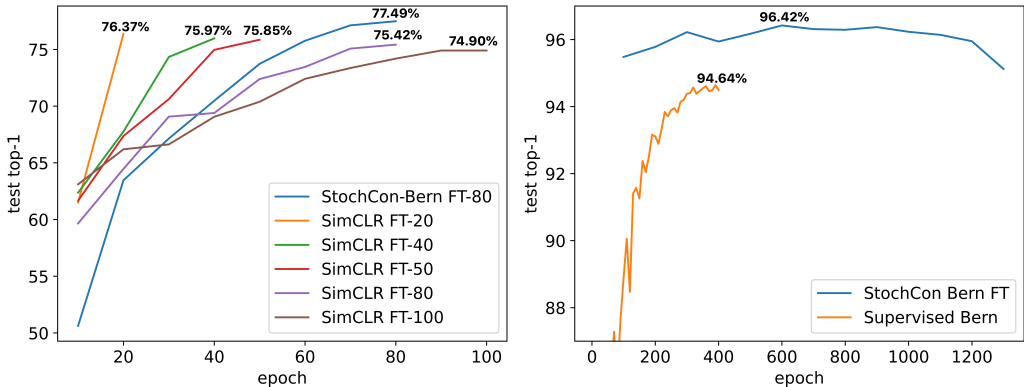


Figure 5: **Left:** Test top-1 performance of multiple trials of SimCLR finetuning on ImageNet. **Right:** Adding a Gumbel Bernoulli to the representation layer ( with  $p = 0.5$  ) of a standard ResNet-50 and training end-to-end on CIFAR10.

To validate that the performance difference in Table 1 was not purely from the finetuning process we perform two experiments:

- ImageNet finetuning:** In Figure 5-Left we finetune multiple SimCLR models over various epoch intervals: {20, 40, 60, 80, 100} and observe that SimCLR does not exceed the reported 76.37% top-1 reported in Table 1 for ImageNet.
- CIFAR10 Bernoulli:** In Figure 5-Right we add a Gumbel Bernoulli layer to the final layer of a standard ResNet50 model (after spatial pooling) and train the model in a standard supervised setting, dropping out the Gumbel-Bernoulli layer with  $p = 0.5$ . The dropout of the layer functions as a proxy to the branch mechanism used in StochCon. We present the best performing model<sup>3</sup>) and note that StochCon outperforms the baseline by 1.78%.

<sup>3</sup>Note that supervised learning typically does not benefit from longer training durations without the use of strong augmentations (Cubuk et al., 2018, 2020).



### A.3 Finetuning procedure

To finetune StochCon, we retain the pre-trained backbone and latent variable distribution, and finetune with Adam (Kingma & Ba, 2015). The *Finetuned* model updates the parameters of the entire network (including the backbone and newly attached linear head), while the *Frozen* model only updates the added linear projection head. We use a learning rate of  $3e-4$ , coupled with a simple step scheduler that scales the learning rate by 0.1 at 80% of training. All our models (including baselines) are trained for various epoch ranges using standard ImageNet augmentations (random flip, random-resized crop), and we report the best performing model in Table 1.

For the Bernoulli-StochCon model, we set the temperature to 0.1 for the entire finetuning process, while the Isotropic-Gaussian distribution only uses the mean (similar to Variational Autoencoders (Kingma & Welling, 2014) at inference time). We suspect that the performance of the *Frozen* Bernoulli-StochCon will match the *Frozen* Isotropic-Gaussian-StochCon model with a properly tuned Gumbel-Bernoulli temperature schedule (Jang et al., 2017; Maddison et al., 2017), but leave this for future work.