

---

# Contrastive Learning Inverts the Data Generating Process

---

Roland S. Zimmermann<sup>\*12</sup> Yash Sharma<sup>\*12</sup> Steffen Schneider<sup>\*123</sup> Matthias Bethge<sup>†1</sup> Wieland Brendel<sup>†1</sup>

## Abstract

Contrastive learning has recently seen tremendous success in self-supervised learning. So far, however, it is largely unclear why the learned representations generalize so effectively to a large variety of downstream tasks. We here prove that feed-forward models trained with objectives belonging to the commonly used InfoNCE family learn to implicitly invert the underlying generative model of the observed data. While the proofs make certain statistical assumptions about the generative model, we observe empirically that our findings hold even if these assumptions are severely violated. Our theory highlights a fundamental connection between contrastive learning, generative modeling, and nonlinear independent component analysis, thereby furthering our understanding of the learned representations as well as providing a theoretical foundation to derive more effective contrastive losses.<sup>1</sup>

## 1. Introduction

With the availability of large collections of unlabeled data, recent work has led to significant advances in self-supervised learning. In particular, contrastive methods have been tremendously successful in learning representations for visual and sequential data (Logeswaran & Lee, 2018; Wu et al., 2018; Oord et al., 2018; Hénaff, 2020; Tian et al., 2019; Hjelm et al., 2019; Bachman et al., 2019; He et al., 2020a; Chen et al., 2020a; Schneider et al., 2019; Baevski et al., 2020a;b; Ravanelli et al., 2020). While a number of explanations have been provided as to why contrastive learning leads to such informative representations, existing theoretical predictions and empirical observations appear to be at odds with each other (Tian et al., 2019; Bachman et al., 2019; Wu et al., 2020; Saunshi et al., 2019).

<sup>\*</sup>Equal contribution <sup>1</sup>University of Tübingen, Tübingen, Germany <sup>2</sup>IMPRS for Intelligent Systems, Tübingen, Germany <sup>3</sup>EPFL, Geneva, Switzerland. Correspondence to: Roland S. Zimmermann <roland.zimmermann@uni-tuebingen.de>.

In a nutshell, contrastive methods aim to learn representations where related samples are aligned (positive pairs, e.g. augmentations of the same image), while unrelated samples are separated (negative pairs) (Chen et al., 2020a). Intuitively, this leads to invariance to irrelevant details or transformations (by decreasing the distance between positive pairs), while preserving a sufficient amount of information about the input for solving downstream tasks (by increasing the distance between negative pairs) (Tian et al., 2020). This intuition has recently been made more precise by (Wang & Isola, 2020), showing that a commonly used contrastive loss from the InfoNCE family (Gutmann & Hyvärinen, 2012; Oord et al., 2018; Chen et al., 2020a) asymptotically converges to a sum of two losses: an *alignment* loss that pulls together the representations of positive pairs, and a *uniformity* loss that maximizes the entropy of the learned latent distribution.

We show that an encoder learned with a contrastive loss from the InfoNCE family can recover the true generative factors of variation (up to rotations) if the process that generated the data fulfills a few weak statistical assumptions. This theory bridges the gap between contrastive learning, nonlinear independent component analysis (ICA) and generative modeling (see Fig. 1). Our theory reveals implicit assumptions encoded in the InfoNCE objective about the generative process underlying the data. If these assumptions are violated, we show a principled way of deriving alternative contrastive objectives based on assumptions regarding the positive pair distribution. We verify our theoretical findings with controlled experiments, providing evidence that our theory holds true in practice, even if the assumptions on the ground-truth generative model are partially violated.

To the best of our knowledge, our work is the first to analyze under what circumstances representation learning methods used in practice provably represent the data in terms of its underlying factors of variation. Our theoretical and empirical results suggest that the success of contrastive learning in many practical applications is due to an implicit and approximate inversion of the data generating process, which explains why the learned representations are useful in a wide range of downstream tasks.

In summary, our contributions are:

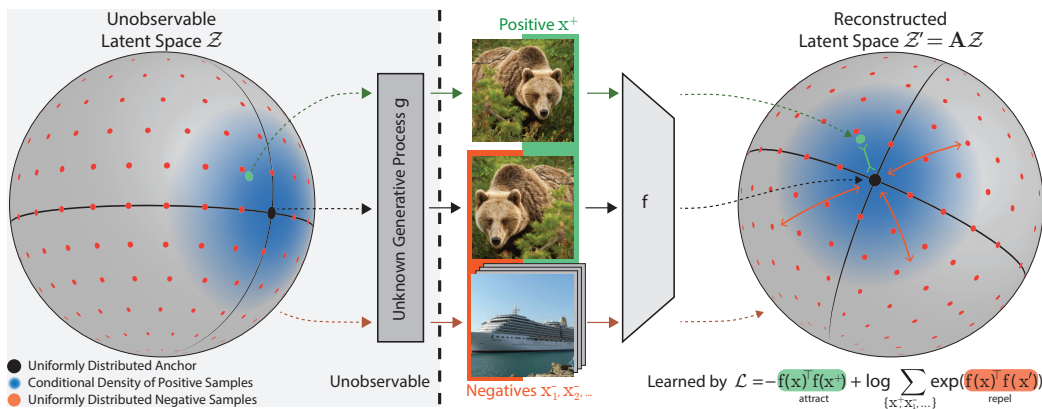


Figure 1. We analyze the setup of contrastive learning, in which a feature encoder  $f$  is trained with the InfoNCE objective (Gutmann & Hyvärinen, 2012; Oord et al., 2018; Chen et al., 2020a) using positive samples (green) and negative samples (orange). We assume the observations are generated by an (unknown) injective generative model  $g$  that maps unobservable latent variables from a hypersphere to observations in another manifold. Under these assumptions, the feature encoder  $f$  implicitly learns to invert the ground-truth generative process  $g$  up to linear transformations, i.e.,  $f = \mathbf{A}g^{-1}$  with an orthogonal matrix  $\mathbf{A}$ , if  $f$  minimizes the InfoNCE objective.

- We establish a theoretical connection between the InfoNCE family of objectives, which is commonly used in self-supervised learning, and nonlinear ICA. We show that training with InfoNCE inverts the data generating process if certain statistical assumptions on the data generating process hold.
- We empirically verify our predictions when the assumed theoretical conditions are fulfilled. In addition, we show successful inversion of the data generating process even if these theoretical assumptions are partially violated.
- We build on top of the CLEVR rendering pipeline (Johnson et al., 2017b) to generate a more visually complex disentanglement benchmark, called *3DIdent*, that contains hallmarks of natural environments (shadows, different lighting conditions, a 3D object, etc.). We demonstrate that a contrastive loss derived from our theoretical framework can identify the ground-truth factors of such complex, high-resolution images.

## 2. Related Work

**Contrastive Learning** Despite the success of contrastive learning (CL), our understanding of the learned representations remains limited, as existing theoretical explanations yield partially contradictory predictions. One way to theoretically motivate CL is to refer to the InfoMax principle (Linsker, 1988), which corresponds to maximizing the mutual information (MI) between different views (Oord et al., 2018; Bachman et al., 2019; Hjelm et al., 2019; Chen et al., 2020a; Tian et al., 2020). However, as optimizing a tighter bound on the MI can produce worse representations (Tschan-

nen et al., 2020), it is not clear how accurate this motivation describes the behavior of CL.

Another approach aims to explain the success by introducing latent classes (Saunshi et al., 2019). While this theory has some appeal, there exists a gap between empirical observations and its predictions, e.g. the prediction that an excessive number of negative samples decreases performance does not corroborate with empirical results (Wu et al., 2018; Tian et al., 2019; He et al., 2020a; Chen et al., 2020a). However, recent work has suggested some empirical evidence for said theoretical prediction, namely, issues with the commonly used sampling strategy for negative samples, and have proposed ways to mitigate said issues as well (Robinson et al., 2020; Chuang et al., 2020).

More recently, the behavior of CL has been analyzed from the perspective of *alignment* and *uniformity* properties of representations, demonstrating that these two properties are correlated with downstream performance (Wang & Isola, 2020). We build on these results to make a connection to cross-entropy minimization from which we can derive identifiability results.

**Nonlinear ICA** Independent Components Analysis (ICA) attempts to find the underlying sources for multidimensional data. In the nonlinear case, said sources correspond to a well-defined nonlinear generative model  $g$ , which is assumed to be invertible (i.e., injective) (Hyvärinen et al., 2001; Jutten et al., 2010). In other words, nonlinear ICA solves a demixing problem: Given observed data  $\mathbf{x} = g(\mathbf{z})$ , it aims to find a model  $f$  that equals the inverse generative model  $g^{-1}$ , which allows for the original sources  $\mathbf{z}$  to be recovered.

Hyvärinen et al. (2019) show that the nonlinear demixing problem can be solved as long as the independent compo-

nents are conditionally mutually independent with respect to some auxiliary variable. The authors further provide practical estimation methods for solving the nonlinear ICA problem (Hyvärinen & Morioka, 2016; 2017), similar in spirit to noise contrastive estimation (NCE; Gutmann & Hyvärinen, 2012). Recent work has generalized this contribution to VAEs (Khemakhem et al., 2020a; Locatello et al., 2020; Klindt et al., 2021), as well as (invertible-by-construction) energy-based models (Khemakhem et al., 2020b). We here extend this line of work to more general feed-forward networks trained using InfoNCE (Oord et al., 2018).

In a similar vein, Roeder et al. (2020) build on the work of Hyvärinen et al. (2019) to show that for a model family which includes InfoNCE, distribution matching implies parameter matching. In contrast, we associate the learned latent representation with the ground-truth generative factors, showing under what conditions the data generating process is inverted, and thus, the true latent factors are recovered.

### 3. Theory

We will show a connection between contrastive learning and identifiability in the form of nonlinear ICA. For this, we introduce a feature encoder  $f$  that maps observations  $\mathbf{x}$  to representations. We consider the widely used *InfoNCE* loss, which often assumes  $L^2$  normalized representations (Wu et al., 2018; He et al., 2020b; Tian et al., 2019; Bachman et al., 2019; Chen et al., 2020a),

$$\mathcal{L}_{\text{contr}}(f; \tau, M) := \mathbb{E}_{\substack{(\mathbf{x}, \tilde{\mathbf{x}}) \sim p_{\text{pos}} \\ \{\mathbf{x}_i^-\}_{i=1}^M \stackrel{\text{i.i.d.}}{\sim} p_{\text{data}}}} \left[ -\log \frac{e^{f(\mathbf{x})^\top f(\tilde{\mathbf{x}})/\tau}}{e^{f(\mathbf{x})^\top f(\tilde{\mathbf{x}})/\tau} + \sum_{i=1}^M e^{f(\mathbf{x}_i^-)^\top f(\tilde{\mathbf{x}})/\tau}} \right]. \quad (1)$$

Here  $M \in \mathbb{Z}_+$  is a fixed number of negative samples,  $p_{\text{data}}$  is the distribution of all observations and  $p_{\text{pos}}$  is the distribution of positive pairs. This loss was motivated by the InfoMax principle (Linsker, 1988), and has been shown to be effective by many recent representation learning methods (Logeswaran & Lee, 2018; Wu et al., 2018; Tian et al., 2019; He et al., 2020a; Hjelm et al., 2019; Bachman et al., 2019; Chen et al., 2020a; Baeviski et al., 2020b). Our theoretical results also hold for a loss function whose denominator only consists of the second summand across the negative samples (e.g., the SimCLR loss (Chen et al., 2020a)).

In the spirit of existing literature on nonlinear ICA (Hyvärinen & Pajunen, 1999; Harmeling et al., 2003; Sprekeler et al., 2014; Hyvärinen & Morioka, 2016; 2017; Gutmann & Hyvärinen, 2012; Hyvärinen et al., 2019; Khemakhem et al., 2020a), we assume that the observations  $\mathbf{x} \in \mathcal{X}$  are generated by an invertible (i.e., injective) generative process  $g: \mathcal{Z} \rightarrow \mathcal{X}$ , where  $\mathcal{X} \subseteq \mathbb{R}^K$  is the space of

observations and  $\mathcal{Z} \subseteq \mathbb{R}^N$  with  $N \leq K$  denotes the space of latent factors. Influenced by the commonly used feature normalization in InfoNCE, we further assume that  $\mathcal{Z}$  is the unit hypersphere  $\mathbb{S}^{N-1}$  (see Appx. A.1.1). Additionally, we assume that the ground-truth marginal distribution of the latents of the generative process is uniform and that the conditional distribution (under which positive pairs have high density) is a von Mises-Fisher (vMF) distribution:

$$p(\mathbf{z}) = |\mathcal{Z}|^{-1}, \quad p(\mathbf{z}|\tilde{\mathbf{z}}) = C_p^{-1} e^{\kappa \mathbf{z}^\top \tilde{\mathbf{z}}} \quad \text{with} \quad (2)$$

$$C_p := \int e^{\kappa \mathbf{z}^\top \tilde{\mathbf{z}}} d\tilde{\mathbf{z}} = \text{const.}, \quad \mathbf{x} = g(\mathbf{z}), \quad \tilde{\mathbf{x}} = g(\tilde{\mathbf{z}}).$$

Given these assumptions, we will show that if  $f$  minimizes the contrastive loss  $\mathcal{L}_{\text{contr}}$ , then  $f$  solves the demixing problem, i.e., inverts  $g$  up to orthogonal linear transformations.

Our theoretical approach consists of three steps: (1) We demonstrate that  $\mathcal{L}_{\text{contr}}$  can be interpreted as the cross-entropy between the (conditional) ground-truth and inferred latent distribution. (2) Next, we show that encoders minimizing  $\mathcal{L}_{\text{contr}}$  maintain distance, i.e., two latent vectors with distance  $\alpha$  in the ground-truth generative model are mapped to points with the same distance  $\alpha$  in the inferred representation. (3) Finally, we leverage distance preservation to show that minimizers of  $\mathcal{L}_{\text{contr}}$  invert the generative process up to orthogonal transformations. Detailed proofs are given in Appx. A.1.2.

Additionally, we will present similar results for general convex bodies in  $\mathbb{R}^N$  and more general similarity measures, see Sec. 3.3. For this, the detailed proofs are given in Appx. A.2.

#### 3.1. Contrastive learning is related to cross-entropy minimization

From the perspective of nonlinear ICA, we are interested in understanding how the representations  $f(\mathbf{x})$  which minimize the contrastive loss  $\mathcal{L}_{\text{contr}}$  (defined in Eq. (1)) are related to the ground-truth source signals  $\mathbf{z}$ . To study this relationship, we focus on the map  $h = f \circ g$  between the recovered source signals  $h(\mathbf{z})$  and the true source signals  $\mathbf{z}$ . Note that this is merely for mathematical convenience; it does not necessitate knowledge regarding neither  $g$  nor the ground-truth factors during learning (beyond the assumptions stated in the theorems).

A core insight is a connection between the contrastive loss and the cross-entropy between the ground-truth latent distribution and a certain model distribution. For this, we expand the theoretical results obtained by Wang & Isola (2020):

**Theorem 1** ( $\mathcal{L}_{\text{contr}}$  converges to the cross-entropy between latent distributions). *If the ground-truth marginal distribution  $p$  is uniform, then for fixed  $\tau > 0$ , as the number of negative samples  $M \rightarrow \infty$ , the (normalized) contrastive*

loss converges to

$$\lim_{M \rightarrow \infty} \mathcal{L}_{\text{contr}}(f; \tau, M) - \log M + \log |\mathcal{Z}| = \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [H(p(\cdot|\mathbf{z}), q_h(\cdot|\mathbf{z}))] \quad (3)$$

where  $H$  is the cross-entropy between the ground-truth conditional distribution  $p$  over positive pairs and a conditional distribution  $q_h$  parameterized by the model  $f$ ,

$$q_h(\tilde{\mathbf{z}}|\mathbf{z}) = C_h(\mathbf{z})^{-1} e^{h(\tilde{\mathbf{z}})^\top h(\mathbf{z})/\tau} \quad (4)$$

with  $C_h(\mathbf{z}) := \int e^{h(\tilde{\mathbf{z}})^\top h(\mathbf{z})/\tau} d\tilde{\mathbf{z}}$ ,

where  $C_h(\mathbf{z}) \in \mathbb{R}^+$  is the partition function of  $q_h$  (see Appx. A.1.1).

Next, we show that the minimizers  $h^*$  of the cross-entropy (4) are isometries in the sense that  $\kappa \mathbf{z}^\top \tilde{\mathbf{z}} = h^*(\mathbf{z})^\top h^*(\tilde{\mathbf{z}})$  for all  $\mathbf{z}$  and  $\tilde{\mathbf{z}}$ . In other words, they preserve the dot product between  $\mathbf{z}$  and  $\tilde{\mathbf{z}}$ .

**Proposition 1** (Minimizers of the cross-entropy maintain the dot product). *Let  $\mathcal{Z} = \mathbb{S}^{N-1}$ ,  $\tau > 0$  and consider the ground-truth conditional distribution of the form  $p(\tilde{\mathbf{z}}|\mathbf{z}) = C_p^{-1} \exp(\kappa \tilde{\mathbf{z}}^\top \mathbf{z})$ . Let  $h$  map onto a hypersphere with radius  $\sqrt{\tau \kappa}$ .<sup>2</sup> Consider the conditional distribution  $q_h$  parameterized by the model, as defined above in Theorem 1, where the hypothesis class for  $h$  (and thus  $f$ ) is assumed to be sufficiently flexible such that  $p(\tilde{\mathbf{z}}|\mathbf{z})$  and  $q_h(\tilde{\mathbf{z}}|\mathbf{z})$  can match. If  $h$  is a minimizer of the cross-entropy  $\mathbb{E}_{p(\tilde{\mathbf{z}}|\mathbf{z})} [-\log q_h(\tilde{\mathbf{z}}|\mathbf{z})]$ , then  $p(\tilde{\mathbf{z}}|\mathbf{z}) = q_h(\tilde{\mathbf{z}}|\mathbf{z})$  and  $\forall \mathbf{z}, \tilde{\mathbf{z}} : \kappa \mathbf{z}^\top \tilde{\mathbf{z}} = h(\mathbf{z})^\top h(\tilde{\mathbf{z}})$ .*

### 3.2. Contrastive learning identifies ground-truth factors on the hypersphere

From the strong geometric property of isometry, we can now deduce a key property of the minimizers  $h^*$ :

**Proposition 2** (Extension of the Mazur-Ulam theorem to hyperspheres and the dot product). *Let  $\mathcal{Z} = \mathbb{S}^{N-1}$ . If  $h : \mathcal{Z} \rightarrow \mathcal{Z}$  maintains the dot product up to a constant factor, i.e.,  $\forall \mathbf{z}, \tilde{\mathbf{z}} \in \mathcal{Z} : \kappa \mathbf{z}^\top \tilde{\mathbf{z}} = h(\mathbf{z})^\top h(\tilde{\mathbf{z}})$ , then  $h$  is an orthogonal linear transformation.*

In the last step, we combine the previous propositions to derive our main result: the minimizers of the contrastive loss  $\mathcal{L}_{\text{contr}}$  solve the demixing problem of nonlinear ICA up to linear transformations, i.e., they identify the original sources  $\mathbf{z}$  for observations  $g(\mathbf{z})$  up to orthogonal linear transformations. For a hyperspherical space  $\mathcal{Z}$  these correspond to combinations of permutations, rotations and sign flips.

<sup>2</sup>Note that in practice this can be implemented as a learnable rescaling operation as the last operation of the network  $f$ .

**Theorem 2.** *Let  $\mathcal{Z} = \mathbb{S}^{N-1}$ , the ground-truth marginal be uniform, and the conditional a vMF distribution (cf. Eq. 2). Let the mixing function  $g$  be differentiable and injective. If the assumed form of  $q_h$ , as defined above, matches that of  $p$ , i.e., both are based on the same metric, and if  $f$  is differentiable and minimizes the CL loss as defined in Eq. (1), then for fixed  $\tau > 0$  and  $M \rightarrow \infty$ ,  $h = f \circ g$  is linear, i.e.,  $f$  recovers the latent sources up to orthogonal linear transformations.*

Note that we do not assume knowledge of the ground-truth generative model  $g$ ; we only make assumptions about the conditional and marginal distribution of the latents. On real data, it is unlikely that the assumed model distribution  $q_h$  can exactly match the ground-truth conditional. We do, however, provide empirical evidence that  $h$  is still an affine transformation even if there is a severe mismatch, see Sec. 4.

### 3.3. Contrastive learning identifies ground-truth factors on convex bodies in $\mathbb{R}^N$

While the previous theoretical results require  $\mathcal{Z}$  to be a hypersphere, we will now show a similar theorem for the more general case of  $\mathcal{Z}$  being a convex body in  $\mathbb{R}^N$ . Note that the hyperrectangle  $[a_1, b_1] \times \dots \times [a_N, b_N]$  is an example of such a convex body.

We follow a similar three step proof strategy as for the hyperspherical case before: (1) We begin again by showing that a properly chosen contrastive loss on convex bodies corresponds to the cross-entropy between the ground-truth conditional and a distribution parametrized by the encoder. For this step, we additionally extend the results of Wang & Isola (2020) to this latent space and loss function. (2) Next, we derive that minimizers of the loss function are isometries of the latent space. Importantly, we do not limit ourselves to a specific metric, thus the result is applicable to a family of contrastive objectives. (3) Finally, we show that these minimizers must be affine transformations. For a special family of conditional distributions (rotationally asymmetric generalized normal distributions (Subbotin, 1923)), we can further narrow the class of solutions to permutations and sign-flips. For the detailed proofs, see Appx. A.2.

As earlier, we assume that the ground-truth marginal distribution of the latents is uniform. However, we now assume that the conditional distribution is exponential:

$$p(\mathbf{z}) = |\mathcal{Z}|^{-1}, \quad p(\mathbf{z}|\tilde{\mathbf{z}}) = C_p^{-1} e^{-\delta(\mathbf{z}, \tilde{\mathbf{z}})} \quad \text{with} \quad (5)$$

$$C_p(\mathbf{z}) := \int e^{-\delta(\mathbf{z}, \tilde{\mathbf{z}})} d\tilde{\mathbf{z}}, \quad \mathbf{x} = g(\mathbf{z}), \quad \tilde{\mathbf{x}} = g(\tilde{\mathbf{z}}),$$

where  $\delta$  is a metric induced by a norm (see Appx. A.2.1).

To reflect the differences between this conditional distribution and the one assumed for the hyperspherical case, we need to introduce an adjusted version of the contrastive loss

in (1):

**Definition 1** ( $\mathcal{L}_{\delta\text{-contr}}$  objective). Let  $\delta : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$  be a metric on  $\mathcal{Z}$ . We define the general InfoNCE loss, which uses  $\delta$  as a similarity measure, as

$$\mathcal{L}_{\delta\text{-contr}}(f; \tau, M) := \mathbb{E}_{\substack{(\mathbf{x}, \tilde{\mathbf{x}}) \sim p_{\text{pos}} \\ \{\mathbf{x}_i\}_{i=1}^M \stackrel{\text{i.i.d.}}{\sim} p_{\text{data}}}} \left[ -\log \frac{e^{-\delta(f(\mathbf{x}), f(\tilde{\mathbf{x}}))/\tau}}{e^{-\delta(f(\mathbf{x}), f(\tilde{\mathbf{x}}))/\tau} + \sum_{i=1}^M e^{-\delta(f(\mathbf{x}_i), f(\tilde{\mathbf{x}}))/\tau}} \right]. \quad (6)$$

Note that this is a generalization of the InfoNCE criterion in Eq. (1). In contrast to the objective above, the representations are no longer assumed to be  $L^2$  normalized, and the dot-product is replaced with a more general similarity measure  $\delta$ .

Analogous to the previously demonstrated case for the hypersphere, for convex bodies  $\mathcal{Z}$ , minimizers of the adjusted  $\mathcal{L}_{\delta\text{-contr}}$  objective solve the demixing problem of nonlinear ICA up to invertible linear transformations:

**Theorem 5.** Let  $\mathcal{Z}$  be a convex body in  $\mathbb{R}^N$ ,  $h = f \circ g : \mathcal{Z} \rightarrow \mathcal{Z}$ , and  $\delta$  be a metric or a semi-metric (cf. Lemma 1 in Appx. A.2.4), induced by a norm. Further, let the ground-truth marginal distribution be uniform and the conditional distribution be as Eq. (5). Let the mixing function  $g$  be differentiable and injective. If the assumed form of  $q_{\text{h}}$  matches that of  $p$ , i.e.,

$$q_{\text{h}}(\tilde{\mathbf{z}}|\mathbf{z}) = C_q^{-1}(\mathbf{z}) e^{-\delta(h(\tilde{\mathbf{z}}), h(\mathbf{z}))/\tau} \\ \text{with } C_q(\mathbf{z}) := \int e^{-\delta(h(\tilde{\mathbf{z}}), h(\mathbf{z}))/\tau} d\tilde{\mathbf{z}}, \quad (7)$$

and if  $f$  is differentiable and minimizes the  $\mathcal{L}_{\delta\text{-contr}}$  objective in Eq. (6) for  $M \rightarrow \infty$ , we find that  $h = f \circ g$  is invertible and affine, i.e., we recover the latent sources up to affine transformations.

Note that the model distribution  $q_{\text{h}}$ , which is implicitly described by the choice of the objective, must be of the same form as the ground-truth distribution  $p$ , i.e., both must be based on the same metric. Thus, identifying different ground-truth conditional distributions requires different contrastive  $\mathcal{L}_{\delta\text{-contr}}$  objectives. This result can be seen as a generalized version of Theorem 2, as it is valid for any convex body  $\mathcal{Z} \subseteq \mathbb{R}^N$ , allowing for a larger variety of conditional distributions.

Finally, under the mild restriction that the ground-truth conditional distribution is based on an  $L^p$  similarity measure for  $p \geq 1, p \neq 2$ ,  $h$  identifies the ground-truth generative factors up to generalized permutations. A generalized permutation matrix  $\mathbf{A}$  is a combination of a permutation and element-wise sign-flips, i.e.,  $\forall \mathbf{z} : (\mathbf{Az})_i = \alpha_i \mathbf{z}_{\sigma(i)}$  with  $\alpha_i = \pm 1$  and  $\sigma$  being a permutation.

**Theorem 6.** Let  $\mathcal{Z}$  be a convex body in  $\mathbb{R}^N$ ,  $h : \mathcal{Z} \rightarrow \mathcal{Z}$ , and  $\delta$  be an  $L^\alpha$  metric or semi-metric (cf. Lemma 1 in Appx. A.2.4) for  $\alpha \geq 1, \alpha \neq 2$ . Further, let the ground-truth marginal distribution be uniform and the conditional distribution be as Eq. (5), and let the mixing function  $g$  be differentiable and invertible. If the assumed form of  $q_{\text{h}}(\cdot|\mathbf{z})$  matches that of  $p(\cdot|\mathbf{z})$ , i.e., both use the same metric  $\delta$  up to a constant scaling factor, and if  $f$  is differentiable and minimizes the  $\mathcal{L}_{\delta\text{-contr}}$  objective in Eq. (6) for  $M \rightarrow \infty$ , we find that  $h = f \circ g$  is a composition of input independent permutations, sign flips and rescaling.

## 4. Experiments

### 4.1. Validation of theoretical claim

We validate our theoretical claims under both perfectly matching and violated conditions regarding the ground-truth marginal and conditional distributions. We consider source signals of dimensionality  $N = 10$ , and sample pairs of source signals in two steps: First, we sample from the marginal  $p(\mathbf{z})$ . For this, we consider both uniform distributions which match our assumptions and non-uniform distributions (e.g., a normal distribution) which violate them. Second, we generate the positive pair by sampling from a conditional distribution  $p(\tilde{\mathbf{z}}|\mathbf{z})$ . Here, we consider matches with our assumptions on the conditional distribution (von Mises-Fisher for  $\mathcal{Z} = \mathbb{S}^{N-1}$ ) as well as violations (e.g. normal, Laplace or generalized normal distribution for  $\mathcal{Z} = \mathbb{S}^{N-1}$ ). Further, we consider spaces beyond the hypersphere, such as the bounded box (which is a convex body) and the unbounded  $\mathbb{R}^N$ .

We generate the observations with a multi-layer perceptron (MLP), following previous work (Hyvärinen & Morioka, 2016; 2017). Specifically, we use three hidden layers with leaky ReLU units and random weights; to ensure that the MLP  $g$  is invertible, we control the condition number of the weight matrices. For our feature encoder  $f$ , we also use an MLP with leaky ReLU units, where the assumed space is denoted by the normalization, or lack thereof, of the encoding. Namely, for the hypersphere (denoted as *Sphere*) and the hyperrectangle (denoted as *Box*) we apply an  $L^2$  and  $L^\infty$  normalization, respectively. For flexibility in practice, we parameterize the normalization magnitude of the *Box*, including it as part of the encoder’s learnable parameters. On the hypersphere we optimize  $\mathcal{L}_{\text{contr}}$  and on the hyperrectangle as well as the unbounded space we optimize  $\mathcal{L}_{\delta\text{-contr}}$ . For further details, see Appx. A.3.

To test for identifiability up to affine transformations, we fit a linear regression between the ground-truth and recovered sources and report the coefficient of determination ( $R^2$ ). To test for identifiability up to generalized permutations, we leverage the mean correlation coefficient (MCC), as used

in previous work (Hyvärinen & Morioka, 2016; 2017). For further details, see Appx. A.3.

We evaluate both identifiability metrics for three different model types. First, we ensure that the problem requires nonlinear demixing by considering the identity function for model  $f$ , which amounts to scoring the observations against the sources (**Identity Model**). Second, we ensure that the problem is solvable within our model class by training our model  $f$  with supervision, minimizing the mean-squared error between  $f(g(\mathbf{z}))$  and  $\mathbf{z}$  (**Supervised Model**). Third, we fit our model without supervision using a contrastive loss (**Unsupervised Model**).

Tables 1 and 2 show results evaluating identifiability up to affine transformations and generalized permutations, respectively. When assumptions match (see column M.), CL recovers a score close to the empirical upper bound. Mismatches in assumptions on the marginal and conditional do not lead to a significant drop in performance with respect to affine identifiability, but do for permutation identifiability compared to the empirical upper bound. In many practical scenarios, we use the learned representations to solve a downstream task, thus, identifiability up to affine transformations is often sufficient. However, for applications where identification of the individual generative factors is desirable, some knowledge of the underlying generative process is required to choose an appropriate loss function and feature normalization. Interestingly, we find that for convex bodies, we obtain identifiability up to permutation even in the case of a normal conditional, which likely is due to the axis-aligned box geometry of the latent domain. Finally, note that the drop in performance for identifiability up to permutations in the last group of Tab. 2 is a natural consequence of either the ground-truth or the assumed conditional being rotationally symmetric, e.g., a normal distribution, in an unbounded space. Here, rotated versions of the latent space are indistinguishable and, thus, the model cannot align the axes of the reconstruction with that of the ground-truth latent space, resulting in a lower score.

To zoom in on how violations of the uniform marginal assumption influence the identifiability achieved by a model in practice, we perform an ablation on the marginal distribution by interpolating between the theoretically assumed uniform distribution and highly locally concentrated distributions. In particular, we consider two cases: (1) a sphere ( $S^9$ ) with a vMF marginal around its north pole for different concentration parameters  $\kappa$ ; (2) a box ( $[0, 1]^{10}$ ) with a normal marginal around the box’s center for different standard deviations  $\sigma$ . For both cases, Fig. 2 shows the  $R^2$  score as a function of the concentration  $\kappa$  and  $1/\sigma^2$  respectively (black). As a reference, the concentration of the used conditional distribution is highlighted as a dashed line. In addition, we also display the probability mass (0–100%)




Figure 2. Varying degrees of violation of the uniformity assumption for the marginal distribution. The figure shows the  $R^2$  score measuring identifiability up to linear transformations (black) as well as the difference between the used marginal and assumed uniform distribution in terms of probability mass (blue) as a function of the marginal’s concentration. The black dotted line indicates the concentration of the used conditional distribution.

that needs to be moved for converting the used marginal distribution (i.e., vMF or normal) into the assumed uniform marginal distribution (blue) as an intuitive measure of the mismatch (i.e.,  $\frac{1}{2} \int |p(\mathbf{z}) - p_{\text{uni}}| d\mathbf{z}$ ). While, we observe significant robustness to mismatch, in both cases, we see performance drop drastically once the marginal distribution is more concentrated than the conditional distribution of positive pairs. In such scenarios, positive pairs are indistinguishable from negative pairs.

## 4.2. Extensions to image data

Previous studies have demonstrated that representation learning using contrastive learning scales well to complex natural image data (Chen et al., 2020a;b; Hénaff, 2020). Unfortunately, the true generative factors of natural images are inaccessible, thus we cannot evaluate identifiability scores.

We consider two alternatives. First, we evaluate on the recently proposed benchmark *KITTI Masks* (Klindt et al., 2021), which is composed of segmentation masks of natural videos. Second, we contribute a novel benchmark (*3DIdent*; cf. Fig. 3) which features aspects of natural scenes, e.g. a complex 3D object and different lighting conditions, while still providing access to the continuous ground-truth factors. For further details, see Appx. A.4.1. *3DIdent* is available at [zenodo.org/record/4502485](https://zenodo.org/record/4502485).

### 4.2.1. KITTI MASKS

*KITTI Masks* (Klindt et al., 2021) is composed of pedestrian segmentation masks extracted from an autonomous driving vision benchmark *KITTI-MOTS* (Geiger et al., 2012), with natural shapes and continuous natural transitions. We compare to SlowVAE (Klindt et al., 2021), the state-of-the-art on the considered dataset. In our experiments, we use the same training hyperparameters (for details see Appx. A.3) and (encoder) architecture as Klindt et al. (2021). The positive

**Contrastive Learning Inverts the Data Generating Process**

Table 1. Identifiability up to affine transformations. Mean  $\pm$  standard deviation over 5 random seeds. Note that only the first row corresponds to a setting that matches ( $\checkmark$ ) our theoretical assumptions, while the others show results for violated assumptions ( $\times$ ; see column  $M$ ). Note that the identity score only depends on the ground-truth space and the marginal distribution defined for the generative process, while the supervised score additionally depends on the space assumed by the model.

| Space     | Generative process $g$ |                           | Space     | Model $f$<br>$q_h(\cdot \cdot)$ | M.           | Identity         | $R^2$ Score [%]  |                  |
|-----------|------------------------|---------------------------|-----------|---------------------------------|--------------|------------------|------------------|------------------|
|           | $p(\cdot)$             | $p(\cdot \cdot)$          |           |                                 |              |                  | Supervised       | Unsupervised     |
| Sphere    | Uniform                | vMF( $\kappa=1$ )         | Sphere    | vMF( $\kappa=1$ )               | $\checkmark$ | $66.98 \pm 2.79$ | $99.71 \pm 0.05$ | $99.42 \pm 0.05$ |
| Sphere    | Uniform                | vMF( $\kappa=10$ )        | Sphere    | vMF( $\kappa=1$ )               | $\times$     | —  —             | —  —             | $99.86 \pm 0.01$ |
| Sphere    | Uniform                | Laplace( $\lambda=0.05$ ) | Sphere    | vMF( $\kappa=1$ )               | $\times$     | —  —             | —  —             | $99.91 \pm 0.01$ |
| Sphere    | Uniform                | Normal( $\sigma=0.05$ )   | Sphere    | vMF( $\kappa=1$ )               | $\times$     | —  —             | —  —             | $99.86 \pm 0.00$ |
| Box       | Uniform                | Normal( $\sigma=0.05$ )   | Unbounded | Normal                          | $\times$     | $67.93 \pm 7.40$ | $99.78 \pm 0.06$ | $99.60 \pm 0.02$ |
| Box       | Uniform                | Laplace( $\lambda=0.05$ ) | Unbounded | Normal                          | $\times$     | —  —             | —  —             | $99.64 \pm 0.02$ |
| Box       | Uniform                | Laplace( $\lambda=0.05$ ) | Unbounded | GenNorm( $\beta=3$ )            | $\times$     | —  —             | —  —             | $99.70 \pm 0.02$ |
| Box       | Uniform                | Normal( $\sigma=0.05$ )   | Unbounded | GenNorm( $\beta=3$ )            | $\times$     | —  —             | —  —             | $99.69 \pm 0.02$ |
| Sphere    | Normal( $\sigma=1$ )   | Laplace( $\lambda=0.05$ ) | Sphere    | vMF( $\kappa=1$ )               | $\times$     | $63.37 \pm 2.41$ | $99.70 \pm 0.07$ | $99.02 \pm 0.01$ |
| Sphere    | Normal( $\sigma=1$ )   | Normal( $\sigma=0.05$ )   | Sphere    | vMF( $\kappa=1$ )               | $\times$     | —  —             | —  —             | $99.02 \pm 0.02$ |
| Unbounded | Laplace( $\lambda=1$ ) | Normal( $\sigma=1$ )      | Unbounded | Normal                          | $\times$     | $62.49 \pm 1.65$ | $99.65 \pm 0.04$ | $98.13 \pm 0.14$ |
| Unbounded | Normal( $\sigma=1$ )   | Normal( $\sigma=1$ )      | Unbounded | Normal                          | $\times$     | $63.57 \pm 2.30$ | $99.61 \pm 0.17$ | $98.76 \pm 0.03$ |

Table 2. Identifiability up to generalized permutations, averaged over 5 runs. Note that while Theorem 6 requires the model latent space to be a convex body and  $p(\cdot|\cdot) = q_h(\cdot|\cdot)$ , we find that empirically either is sufficient. The results are grouped in four blocks corresponding to different types and degrees of violation of assumptions of our theory showing identifiability up to permutations: (1) no violation, violation of the assumptions on either the (2) space or (3) the conditional distribution, or (4) both.

| Space | Generative process $g$ |                                    | Space     | Model $f$<br>$q_h(\cdot \cdot)$ | M.           | Identity         | MCC Score [%]    |                  |
|-------|------------------------|------------------------------------|-----------|---------------------------------|--------------|------------------|------------------|------------------|
|       | $p(\cdot)$             | $p(\cdot \cdot)$                   |           |                                 |              |                  | Supervised       | Unsupervised     |
| Box   | Uniform                | Laplace( $\lambda=0.05$ )          | Box       | Laplace                         | $\checkmark$ | $46.55 \pm 1.34$ | $99.93 \pm 0.03$ | $98.62 \pm 0.05$ |
| Box   | Uniform                | GenNorm( $\beta=3; \lambda=0.05$ ) | Box       | GenNorm( $\beta=3$ )            | $\checkmark$ | —  —             | —  —             | $99.90 \pm 0.06$ |
| Box   | Uniform                | Normal( $\sigma=0.05$ )            | Box       | Normal                          | $\times$     | —  —             | —  —             | $99.77 \pm 0.01$ |
| Box   | Uniform                | Laplace( $\lambda=0.05$ )          | Box       | Normal                          | $\times$     | —  —             | —  —             | $99.76 \pm 0.02$ |
| Box   | Uniform                | GenNorm( $\beta=3; \lambda=0.05$ ) | Box       | Laplace                         | $\times$     | —  —             | —  —             | $98.80 \pm 0.02$ |
| Box   | Uniform                | Laplace( $\lambda=0.05$ )          | Unbounded | Laplace                         | $\times$     | —  —             | $99.97 \pm 0.03$ | $98.57 \pm 0.02$ |
| Box   | Uniform                | GenNorm( $\beta=3; \lambda=0.05$ ) | Unbounded | GenNorm( $\beta=3$ )            | $\times$     | —  —             | —  —             | $99.85 \pm 0.01$ |
| Box   | Uniform                | Normal( $\sigma=0.05$ )            | Unbounded | Normal                          | $\times$     | —  —             | —  —             | $58.26 \pm 3.00$ |
| Box   | Uniform                | Laplace( $\lambda=0.05$ )          | Unbounded | Normal                          | $\times$     | —  —             | —  —             | $59.67 \pm 2.33$ |
| Box   | Uniform                | Normal( $\sigma=0.05$ )            | Unbounded | GenNorm( $\beta=3$ )            | $\times$     | —  —             | —  —             | $43.80 \pm 2.15$ |

pairs consist of nearby frames with a time separation  $\overline{\Delta t}$ .

As argued and shown in Klindt et al. (2021), the transitions in the ground-truth latents between nearby frames is sparse. Unsurprisingly then, Table 3 shows that assuming a Laplace conditional as opposed to a normal conditional in the contrastive loss leads to better identification of the underlying factors of variation. SlowVAE also assumes a Laplace conditional (Klindt et al., 2021) but appears to struggle if the frames of a positive pair are too similar ( $\overline{\Delta t} = 0.05s$ ). This degradation in performance is likely due to the limited expressiveness of the decoder deployed in SlowVAE.

#### 4.2.2. 3DIDENT

**Dataset description** We build on (Johnson et al., 2017b) and use the Blender rendering engine (Blender Online Com-

Table 3. **KITTI Masks**. Mean  $\pm$  standard deviation over 10 random seeds.  $\overline{\Delta t}$  indicates the average temporal distance of frames used.

|                               | Model   | Model Space | MCC [%]        |
|-------------------------------|---------|-------------|----------------|
| $\overline{\Delta t} = 0.05s$ | SlowVAE | Unbounded   | $66.1 \pm 4.5$ |
|                               | Laplace | Unbounded   | $77.1 \pm 1.0$ |
|                               | Laplace | Box         | $74.1 \pm 4.4$ |
|                               | Normal  | Unbounded   | $58.3 \pm 5.4$ |
|                               | Normal  | Box         | $59.9 \pm 5.5$ |
| $\overline{\Delta t} = 0.15s$ | SlowVAE | Unbounded   | $79.6 \pm 5.8$ |
|                               | Laplace | Unbounded   | $79.4 \pm 1.9$ |
|                               | Laplace | Box         | $80.9 \pm 3.8$ |
|                               | Normal  | Unbounded   | $60.2 \pm 8.7$ |
|                               | Normal  | Box         | $68.4 \pm 6.7$ |




Figure 3. **3DIdent**. Influence of the latent factors  $\mathbf{z}$  on the renderings  $\mathbf{x}$ . Each column corresponds to a traversal in one of the ten latent dimensions while the other dimensions are kept fixed.

munity, 2021) to create visually complex 3D images (see Fig. 3). Each image in the dataset shows a colored 3D object which is located and rotated above a colored ground in a 3D space. Additionally, each scene contains a colored spotlight focused on the object and located on a half-circle around the scene. The observations are encoded with an RGB color space, and the spatial resolution is  $224 \times 224$  pixels.

The images are rendered based on a 10-dimensional latent, where: (1) three dimensions describe the XYZ position, (2) three dimensions describe the rotation of the object in Euler angles, (3) two dimensions describe the color of the object and the ground of the scene, respectively, and (4) two dimensions describe the position and color of the spotlight. We use the HSV color space to describe the color of the object and the ground with only one latent each by having the latent factor control the hue value. For more details on the dataset see Sec. A.4.

The dataset contains 250 000 observation-latent pairs where the latents are uniformly sampled from the hyperrectangle  $\mathcal{Z}$ . To sample positive pairs  $(\mathbf{z}, \tilde{\mathbf{z}})$  we first sample a value  $\tilde{\mathbf{z}}'$  from the data conditional  $p(\tilde{\mathbf{z}}'|\mathbf{z})$ , and then use nearest-neighbor matching<sup>3</sup> implemented by FAISS (Johnson et al., 2017a) to find the latent  $\tilde{\mathbf{z}}$  closest to  $\tilde{\mathbf{z}}'$  (in  $L^2$  distance) for which there exists an image rendering. In addition, unlike previous work (Locatello et al., 2019), we create a hold-out test set with 25 000 distinct observation-latent pairs.

**Experiments and Results** We train a convolutional feature encoder  $f$  composed of a ResNet18 architecture (He

<sup>3</sup>We used an Inverted File Index (IVF) with Hierarchical Navigable Small World (HNSW) graph exploration for fast indexing.

et al., 2016) and an additional fully-connected layer, with a LeakyReLU nonlinearity as the hidden activation. For more details, see Appx. A.3. Following the same methodology as in Sec. 4.1, i) depending on the assumed space, the output of the feature encoder is normalized accordingly and ii) in addition to the CL models, we also train a supervised model to serve as an upper bound on performance. We consider normal and Laplace distributions for positive pairs. Note, that due to the finite dataset size we only sample from an approximation of these distributions.

As in Tables 1 and 2, the results in Table 4 demonstrate that CL reaches scores close to the topline (supervised) performance, and mismatches between the assumed and ground-truth conditional distribution do not harm the performance significantly. However, if the hypothesis class of the encoder is too restrictive to model the ground-truth conditional distribution, we observe a clear drop in performance, i.e., mapping a box onto a sphere. Note, that this corresponds to the InfoNCE objective for  $L^2$ -normalized representations, commonly used for self-supervised representation learning (Wu et al., 2018; He et al., 2020b; Tian et al., 2019; Bachman et al., 2019; Chen et al., 2020a). Finally, the last result shows that leveraging image augmentations (Chen et al., 2020a) as opposed to sampling from a specified conditional distribution of positive pairs  $p(\cdot|\cdot)$  results in a performance drop. For details on the experiment, see Appx. Sec. A.3. We explain this with the greater mismatch between the conditional distribution assumed by the model and the conditional distribution induced by the augmentations. In all, we demonstrate validation of our theoretical claims even for generative processes with higher visual complexity than those considered in Sec. 4.1.



Table 4. Identifiability up to affine transformations on the test set of 3DIdent. Mean  $\pm$  standard deviation over 3 random seeds. As earlier, only the first row corresponds to a setting that matches the theoretical assumptions for linear identifiability; the others show distinct violations. Supervised training with unbounded space achieves scores of  $R^2 = (98.67 \pm 0.03)\%$  and  $MCC = (99.33 \pm 0.01)\%$ . The last row refers to using the image augmentations suggested by Chen et al. (2020a) to generate positive image pairs. For performance on the training set, see Appx. Table 5.

| Dataset<br>$p(\cdot \cdot)$ | Model $f$ |                    | M. | Identity [%]    | Unsupervised [%] |                  |
|-----------------------------|-----------|--------------------|----|-----------------|------------------|------------------|
|                             | Space     | $q_h(\cdot \cdot)$ |    | $R^2$           | $R^2$            | MCC              |
| Normal                      | Box       | Normal             | ✓  | $5.25 \pm 1.20$ | $96.73 \pm 0.10$ | $98.31 \pm 0.04$ |
| Normal                      | Unbounded | Normal             | ✗  | —  —            | $96.43 \pm 0.03$ | $54.94 \pm 0.02$ |
| Laplace                     | Box       | Normal             | ✗  | —  —            | $96.87 \pm 0.08$ | $98.38 \pm 0.03$ |
| Normal                      | Sphere    | vMF                | ✗  | —  —            | $65.74 \pm 0.01$ | $42.44 \pm 3.27$ |
| Augm.                       | Sphere    | vMF                | ✗  | —  —            | $45.51 \pm 1.43$ | $46.34 \pm 1.59$ |

## 5. Conclusion

We showed that objectives belonging to the InfoNCE family, the basis for a number of state-of-the-art techniques in self-supervised representation learning, can uncover the true generative factors of variation underlying the observational data. To succeed, these objectives implicitly encode a few weak assumptions about the statistical nature of the underlying generative factors. While these assumptions will likely not be exactly matched in practice, we showed empirically that the underlying factors of variation are identified even if theoretical assumptions are severely violated.

Our theoretical and empirical results suggest that the representations found with contrastive learning implicitly (and approximately) invert the generative process of the data. This could explain why the learned representations are so useful in many downstream tasks. It is known that a decisive aspect of contrastive learning is the right choice of augmentations that form a positive pair. We hope that our framework might prove useful for clarifying the ways in which certain augmentations affect the learned representations, and for finding improved augmentation schemes.

Furthermore, our work opens avenues for constructing more effective contrastive losses. As we demonstrate, imposing a contrastive loss informed by characteristics of the latent space can considerably facilitate inferring the correct semantic descriptors, and thus boost performance in downstream tasks. While our framework already allows for a variety of *conditional* distributions, it is an interesting open question how to adapt it to *marginal* distributions beyond the uniform implicitly encoded in InfoNCE. Also, future work may extend our theoretical framework by incorporating additional assumptions about our visual world, such as compositionality, hierarchy or objectness. Accounting for such inductive biases holds enormous promise in forming the basis for the next generation of self-supervised learning algorithms.

Taken together, we lay a strong theoretical foundation for not only understanding but extending the success of state-of-the-art self-supervised learning techniques.

## Author contributions

The project was initiated by WB. RSZ, SS and WB jointly derived the theory. RSZ and YS implemented and executed the experiments. The 3DIdent dataset was created by RSZ with feedback from SS, YS, WB and MB. RSZ, YS, SS and WB contributed to the final version of the manuscript.

## Acknowledgements

We thank Muhammad Waleed Gondal, Ivan Ustyuzhaninov, David Klindt, Lukas Schott and Luisa Eck for helpful discussions. We thank Bozidar Antic, Shubham Krishna and Jugoslav Stojcheski for ideas regarding the design of 3DIdent. We thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting RSZ, YS and StS. StS acknowledges his membership in the European Laboratory for Learning and Intelligent Systems (ELLIS) PhD program. We acknowledge support from the German Federal Ministry of Education and Research (BMBF) through the Competence Center for Machine Learning (TUE.AI, FKZ 01IS18039A) and the Bernstein Computational Neuroscience Program Tübingen (FKZ: 01GQ1002). WB acknowledges support via his Emmy Noether Research Group funded by the German Science Foundation (DFG) under grant no. BR 6382/1-1 as well as support by Open Philanthropy and the Good Ventures Foundation. MB and WB acknowledge funding from the MICrONS program of the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DoI/IBC) contract number D16PC00003.

## References

- Bachman, P., Hjelm, R. D., and Buchwalter, W. Learning representations by maximizing mutual information across views. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 15509–15519, 2019.
- Baevski, A., Schneider, S., and Auli, M. vq-wav2vec: Self-supervised learning of discrete speech representations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020a.
- Baevski, A., Zhou, Y., Mohamed, A., and Auli, M. wav2vec 2.0: A framework for self-supervised learning of speech representations. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020b.
- Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Blender Institute, Amsterdam, 2021.
- Burgess, C. and Kim, H. 3d shapes dataset. <https://github.com/deepmind/3dshapes-dataset/>, 2018.
- Čalka, A. Local isometries of compact metric spaces. *Proceedings of the American Mathematical Society*, 85(4): 643–647, 1982.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. E. 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.
- Chen, T., Kornblith, S., Swersky, K., Norouzi, M., and Hinton, G. E. Big self-supervised models are strong semi-supervised learners. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020b.
- Chuang, C., Robinson, J., Lin, Y., Torralba, A., and Jegelka, S. Debiased contrastive learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Deng, J., Dong, W., Socher, R., Li, L., Li, K., and Li, F. 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.
- Dittadi, A., Träuble, F., Locatello, F., Wüthrich, M., Agrawal, V., Winther, O., Bauer, S., and Schölkopf, B. On the transfer of disentangled representations in realistic settings. *International Conference on Learning Representations (ICLR)*, 2021.
- Geiger, A., Lenz, P., and Urtasun, R. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012*, pp. 3354–3361. IEEE Computer Society, 2012. doi: 10.1109/CVPR.2012.6248074.
- Gondal, M. W., Wuthrich, M., Miladinovic, D., Locatello, F., Breidt, M., Volchkov, V., Akpo, J., Bachem, O., Schölkopf, B., and Bauer, S. On the transfer of inductive bias from simulation to the real world: a new disentanglement dataset. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 15714–15725, 2019.
- Gutmann, M. U. and Hyvärinen, A. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *The Journal of Machine Learning Research*, 13:307–361, 2012.
- Harmeling, S., Ziehe, A., Kawanabe, M., and Müller, K.-R. Kernel-based nonlinear blind source separation. *Neural Computation*, 15(5):1089–1124, 2003.
- He, K., Zhang, X., Ren, S., and Sun, J. 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.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. B. Momentum contrast for unsupervised visual representation learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 9726–9735. IEEE, 2020a. doi: 10.1109/CVPR42600.2020.00975.

- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. B. Momentum contrast for unsupervised visual representation learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 9726–9735. IEEE, 2020b. doi: 10.1109/CVPR42600.2020.00975.
- Hénaff, O. J. Data-efficient image recognition with contrastive predictive coding. 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. 4182–4192. PMLR, 2020.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- Hyvärinen, A. and Morioka, H. Unsupervised feature extraction by time-contrastive learning and nonlinear ICA. In Lee, D. D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3765–3773, 2016.
- Hyvärinen, A. and Morioka, H. Nonlinear ICA of temporally dependent stationary sources. In Singh, A. and Zhu, X. J. (eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, volume 54 of *Proceedings of Machine Learning Research*, pp. 460–469. PMLR, 2017.
- Hyvärinen, A. and Pajunen, P. Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 12(3):429–439, 1999.
- Hyvärinen, A., Karhunen, J., and Oja, E. *Independent Component Analysis*. Wiley Interscience, 2001.
- Hyvärinen, A., Sasaki, H., and Turner, R. E. Nonlinear ICA using auxiliary variables and generalized contrastive learning. In Chaudhuri, K. and Sugiyama, M. (eds.), *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pp. 859–868. PMLR, 2019.
- Johnson, J., Douze, M., and Jégou, H. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017a.
- Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Zitnick, C. L., and Girshick, R. B. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pp. 1988–1997. IEEE Computer Society, 2017b. doi: 10.1109/CVPR.2017.215.
- Jutten, C., Babaie-Zadeh, M., and Karhunen, J. Nonlinear mixtures. *Handbook of Blind Source Separation, Independent Component Analysis and Applications*, pp. 549–592, 2010.
- Khemakhem, I., Kingma, D. P., Monti, R. P., and Hyvärinen, A. Variational autoencoders and nonlinear ICA: A unifying framework. In Chiappa, S. and Calandra, R. (eds.), *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2207–2217. PMLR, 2020a.
- Khemakhem, I., Monti, R. P., Kingma, D. P., and Hyvärinen, A. Ice-beem: Identifiable conditional energy-based deep models based on nonlinear ICA. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020b.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Klindt, D., Schott, L., Sharma, Y., Ustyuzhaninov, I., Brendel, W., Bethge, M., and Paiton, D. Towards nonlinear disentanglement in natural data with temporal sparse coding. *International Conference on Learning Representations (ICLR)*, 2021.
- Lamperti, J. et al. On the isometries of certain function-spaces. *Pacific J. Math*, 8(3):459–466, 1958.
- Lee, J. M. Smooth manifolds. In *Introduction to Smooth Manifolds*, pp. 606–607. Springer, 2013.
- Li, C.-K. and So, W. Isometries of  $\ell_p$ -norm. *The American Mathematical Monthly*, 101(5):452–453, 1994.
- Linsker, R. Self-organization in a perceptual network. *Computer*, 21(3):105–117, 1988.
- Locatello, F., Bauer, S., Lucic, M., Rätsch, G., Gelly, S., Schölkopf, B., and Bachem, O. Challenging common assumptions in the unsupervised learning of disentangled

- representations. In Chaudhuri, K. and Salakhutdinov, R. (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. 4114–4124. PMLR, 2019.
- Locatello, F., Poole, B., Rätsch, G., Schölkopf, B., Bachem, O., and Tschannen, M. Weakly-supervised disentanglement without compromises. 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. 6348–6359. PMLR, 2020.
- Logeswaran, L. and Lee, H. An efficient framework for learning sentence representations. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- Mankiewicz, P. Extension of isometries in normed linear spaces. *Bulletin de l'Academie polonaise des sciences: Serie des sciences mathematiques, astronomiques et physiques*, 20(5):367–+, 1972.
- Newell, M. E. *The Utilization of Procedure Models in Digital Image Synthesis*. PhD thesis, The University of Utah, 1975. AAI7529894.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Ravanelli, M., Zhong, J., Pascual, S., Swietojanski, P., Monteiro, J., Trmal, J., and Bengio, Y. Multi-task self-supervised learning for robust speech recognition. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pp. 6989–6993. IEEE, 2020. doi: 10.1109/ICASSP40776.2020.9053569.
- Robinson, J., Chuang, C.-Y., Sra, S., and Jegelka, S. Contrastive learning with hard negative samples. *arXiv preprint arXiv:2010.04592*, 2020.
- Roeder, G., Metz, L., and Kingma, D. P. On linear identifiability of learned representations. *arXiv preprint arXiv:2007.00810*, 2020.
- Ruzhansky, M. and Sugimoto, M. On global inversion of homogeneous maps. *Bulletin of Mathematical Sciences*, 5(1):13–18, 2015.
- Saunshi, N., Plevrakis, O., Arora, S., Khodak, M., and Khandeparkar, H. A theoretical analysis of contrastive unsupervised representation learning. In Chaudhuri, K. and Salakhutdinov, R. (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. 5628–5637. PMLR, 2019.
- Schneider, S., Baevski, A., Collobert, R., and Auli, M. wav2vec: Unsupervised pre-training for speech recognition. *CoRR*, abs/1904.05862, 2019.
- Sprekeler, H., Zito, T., and Wiskott, L. An extension of slow feature analysis for nonlinear blind source separation. *The Journal of Machine Learning Research*, 15(1):921–947, 2014.
- Subbotin, M. F. On the law of frequency of error. *Mat. Sb.*, 31(2):296–301, 1923.
- Tian, Y., Krishnan, D., and Isola, P. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., and Isola, P. What makes for good views for contrastive learning, 2020.
- Tschannen, M., Djolonga, J., Rubenstein, P. K., Gelly, S., and Lucic, M. On mutual information maximization for representation learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- Wang, T. and Isola, P. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. 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. 9929–9939. PMLR, 2020.
- Wu, M., Zhuang, C., Yamins, D., and Goodman, N. On the importance of views in unsupervised representation learning. 2020.
- Wu, Z., Xiong, Y., Yu, S. X., and Lin, D. Unsupervised feature learning via non-parametric instance discrimination. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 3733–3742. IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00393.