

---

# Position: An Empirically Grounded Identifiability Theory Will Accelerate Self-Supervised Learning Research

---

Patrik Reizinger<sup>1 2</sup> Randall Balestriero<sup>3</sup> David Klindt<sup>4</sup> Wieland Brendel<sup>1 2 5</sup>

## Abstract

Self-Supervised Learning (SSL) powers many current AI systems. As research interest and investment grow, the SSL design space continues to expand. The Platonic view of SSL, following the *Platonic Representation Hypothesis* (PRH), suggests that despite different methods and engineering approaches, all representations converge to the same Platonic ideal. However, this phenomenon lacks precise theoretical explanation. By synthesizing evidence from Identifiability Theory (IT), we show that the PRH can emerge in SSL. There is a gap between SSL theory and practice: Current IT cannot explain SSL’s empirical success, though it has practically relevant insights. Our work formulates a blueprint for SSL research to bridge this gap: we propose expanding IT into what we term Singular Identifiability Theory (SITh), a broader theoretical framework encompassing the entire SSL pipeline. SITh would allow deeper insights into the implicit data assumptions in SSL and advance the field towards learning more interpretable and generalizable representations. We highlight three critical directions for future research: 1) training dynamics and convergence properties of SSL; 2) the impact of finite samples, batch size, and data diversity; and 3) the role of inductive biases in architecture, augmentations, initialization schemes, and optimizers.

## 1. Introduction

Self-supervised learning (SSL) drives many current AI breakthroughs in language, vision, and image-text models (Radford et al., 2018; 2021; Assran et al., 2023). This

remarkable success builds on a plethora of algorithms and engineering practices. However, the model zoo makes SSL hard to navigate, divides the attention for a better theoretical understanding and, we argue, dampens future advancements. SSL is likely to continue being a foundational paradigm of machine learning, but it needs to be consolidated, similar to the program of unifying deep learning through geometry (Bronstein et al., 2021). There were recent attempts to unify the SSL model zoo (Wang & Isola, 2020; Morningstar et al., 2024; Bizeul et al., 2024; Fleissner et al., 2025). Huh et al. (2024) even argued, by positing the Platonic Representation Hypothesis (PRH), that neural networks converge to effectively the same representation, despite different setups, loss functions, data modalities, and engineering practices. However, their analysis does not provide a conclusive mathematical answer *when and why* this happens. We do have theories that can provide such an answer, but a gap exists between SSL theory and practice. Prevalent frameworks such as identifiability theory (IT) cannot explain the role of initialization, stop gradients, or learning dynamics, and neither can solve problems like dimensional collapse (Jing et al., 2022) or loss saturation. An even larger problem of current theory is generalization: The promise of self-supervised representations to work well on many downstream tasks alludes to “universality,” or at least out-of-distribution (OOD) generalization. However, current theories mostly focus on the i.i.d. case, and empirical methods generally perform well only in-distribution (Montero et al., 2021; Schott et al., 2021; Montero et al., 2024; Mayilvahanan et al., 2024a;b).

**Position:** empirical advancements alone are insufficient to accelerate SSL research, we need to bridge the gap between theory and practice.

We need an empirically grounded theory, which we call Singular Identifiability Theory (SITh), to accelerate SSL research by:

- (i) Designing realistic data generating processes (DGPs), grounded in empirical observations;
- (ii) Formalizing when and why self-supervised representations are converging;
- (iii) Providing principled recommendations for designing and evaluating SSL.

<sup>1</sup>Max Planck Institute for Intelligent Systems, Tübingen, Germany <sup>2</sup>ELLIS Institute Tübingen, Germany <sup>3</sup>Department of Computer Science, Brown University, Rhode Island, USA <sup>4</sup>Cold Spring Harbor Laboratory, Cold Spring Harbor, New York, USA <sup>5</sup>Tübingen AI Center, Germany. Correspondence to: Patrik Reizinger <patrik.reizinger@tuebingen.mpg.de>.

**This position paper argues that to accelerate SSL research, we need to develop a theoretical framework grounded in empirical observation.** Acknowledging the success of IT, we call for its extension to what we term Singular Identifiability Theory (SITh)—alluding to Singular Learning Theory (SLT) (Watanabe, 2009; 2020), which extended learning theory (LT) to considerations relevant to neural networks. SSL practitioners need to rely on principled guidance derived from theory, whereas theoreticians need to concentrate on studying practically relevant aspects of SSL—we provide concrete examples and research questions in Tab. 1. Our **contributions** are:

- We provide an approachable introduction to identifiability theory (IT) for practitioners, demonstrating how SSL practice can benefit from IT (§ 3);
- We call for an empirically grounded extension of identifiability, termed Singular Identifiability Theory (SITh) to close the gap between SSL theory and practice, mainly relying on the concept of a DGP (§ 4);
- We provide an extensive overview of the open questions in SSL, formulate concrete research questions and promising research directions to suggest how empirical observations and theoretical results can drive SSL forward (Tab. 1).

## 2. Background

**SSL drives most AI breakthroughs.** Long gone are the days when supervised, semi-supervised, and unsupervised learning co-existed as the main tracks of top AI conferences. The emergence of Internet-scale datasets coupled with the need for foundation models (Bommasani et al., 2021)—models capable of solving many tasks post-training—has moved the landscape of AI under one umbrella: self-supervised learning (SSL). With access to data and compute, hundreds of SSL methods have emerged (Balestriero et al., 2023). As a result, SSL now holds the state-of-the-art across numerous domains like remote sensing and medical data (Wang et al., 2022; Krishnan et al., 2022), and applications such as zero-shot reasoning (Radford et al., 2021; Zhai et al., 2023), multi-modal learning (Girdhar et al., 2022). The distinct engineering choices and implementation details make it hard to empirically test every hypothesis one may have in mind. Yet, such hypothesis testing is becoming more important than ever to understand possible robustness, fairness, compositional, or structural limitations of existing methods. Without a clear understanding of today’s limitations, it is getting harder to develop tomorrow’s methods.

**Self-supervised representations are often but not always similar.** Theoretical and empirical evidence indicates a surprising similarity between some self-supervised representations. Roeder et al. (2020) proved that in specific cases, two trained models learn a similar representation up to a linear representation. Works on relative representations and model stitching (Moschella et al., 2023; Cannistraci et al., 2023; Norelli et al., 2023; Fumero et al., 2024; Maiorca

et al., 2024) demonstrated that learned self-supervised representations can be related via simple (e.g., linear) transformations. A similar notion of representation linearity was demonstrated in language models (Park et al., 2023; Marconato et al., 2024; Jha et al., 2025), building from earlier observations of analogy making in word embeddings like word2vec (Mikolov et al., 2013; Arora et al., 2016). The PRH (Huh et al., 2024) proposed an explanation of why such convergence can happen, alluding to Plato’s allegory of the cave. Huh et al. (2024) formalized three hypotheses, corresponding to the effect of: i) *function class*, ii) *data sets and loss functions*, and iii) *regularization*. Earlier work by Duan et al. (2020) suggested a link between model convergence and identifiability. We discuss evidence for their importance in SSL and highlight that current theories cannot fully explain their effect (cf. Appx. B for details). Morningstar et al. (2024) concluded that the data augmentations matter more than the SSL method for downstream classification accuracy on ImageNet (Krizhevsky et al., 2012). Ciernik et al. (2024) demonstrated that SSL representations are often, but not universally, similar (cf. Fig. 1). The prevalent problems of dimensional collapse (Jing et al., 2022) or the projector phenomenon (Jing et al., 2022; Chen et al., 2020; von Kügelgen et al., 2021), suggest that the PRH can only be true for a set of models. Thus, a mathematically conclusive answer as of *when and why* the PRH can hold (in SSL) eludes the research community. Our work aims to guide research interest towards answering these questions.

## 3. Identifiability theory

This section provides an intuitive introduction to the terminology of identifiability theory (IT) (§ 3.1), and demonstrates IT’s merits for practitioners on the example of SimCLR (§ 3.2). We also review the state of the field (§ 3.3); and highlight many IT-driven applications (§ 3.4).

### 3.1. The terminology of identifiability

Consider a diverse collection of images like ImageNet. For each image, certain “ingredients” or latents came together to create it: the type of object, its position, color, lighting conditions, and so on. Identifiability theory (IT) provides a mathematical framework to study whether we can discover these underlying properties just by looking at the images.

Imagine these images were created by a virtual “renderer”—what IT calls a *data generating process (DGP)*—that takes these latents as input and produces the final image (Kulkarni et al., 2015). For instance, to generate a photo of a red car, this renderer would need inputs specifying “car” as the object class, “red” for its color, and values for its position. The key question IT asks is:

*Given only the images, under what conditions can we work backwards to uncover the original latents through SSL?*

This is not always straightforward since there might be mul-

multiple valid ways to describe the same latent. Take color: we could represent it using RGB values or HSV coordinates—both produce identical images (Higgins et al., 2018). IT handles this by grouping these equivalent descriptions into *equivalence classes*. To pin down a specific solution, we need additional assumptions about the DGP, e.g., regarding the distribution of latents or the function class of the renderer (Sprekeler et al., 2014).

### 3.2. What can identifiability bring to SSL? A gentle introduction with SimCLR

To illustrate how IT can improve our understanding of SSL methods, we consider SimCLR (Chen et al., 2020).

**Step 1: the birth of SimCLR.** SimCLR was proposed from an empirical perspective without any identifiability theoretic considerations. Its success inspired theoreticians to try to understand its operating principles.

**Step 2: Uniformity and alignment.** The first theoretical result by Wang & Isola (2020) conceptualized the SimCLR loss in terms of two opposing “forces:” an attractive force termed *alignment* pulling the representations of similar (augmented, positive) samples together, whereas a repelling force termed *uniformity* pushing dissimilar (negative) samples apart. This conceptual framework provided some practical guidance on design choices, including the temperature value or batch size. However, problems such as dimensional collapse (Jing et al., 2022) (when some latent information is not learned), or the emergence of the projector—when downstream performance is better for the non-ultimate layer—remained unexplained.

**Step 3: Identifiable DGP.** Zimmermann et al. (2021) showed—building on prior works in nonlinear Independent Component Analysis (ICA) (Hyvärinen & Morioka, 2016; Hyvärinen et al., 2019; Khemakhem et al., 2020a,b; Hyvärinen & Morioka, 2017)—that the SimCLR objective is optimized when the neural network learns to parametrize a DGP on the hypersphere, where the augmentations are modeled with a von Mises-Fisher (vMF) conditional distribution. This result showed how a few assumptions are sufficient for identifiability: a hypersphere latent space with uniformly distributed samples, augmentation distributions according to a vMF, and an invertible encoder. Furthermore, this result made the *implicit* assumptions about the data and the underlying DGP explicit and pinpointed their unrealistic nature: the vMF conditional is isotropic, which does not correspond to realistic augmentations, including large crops, which are crucial for SSL. Thus, IT can model *simplified* SSL methods, but not the practically used versions.

**Step 4: towards a more realistic theory.** A later extension by Rusak et al. (2024) proved identifiability for a conditional which can model augmentations with differ-

ent strengths. Yet, we still do not understand (and cannot mitigate) dimensional collapse (Jing et al., 2022) or the projector.

### 3.3. The state of identifiability theory

The prevalent method family with identifiability guarantees is Independent Component Analysis (ICA) (Comon, 1994; Hyvärinen & Pajunen, 1999; Hyvärinen et al., 2001; Hyvärinen et al., 2023). ICA methods aim to extract (conditionally) independent latent variables, called sources—this problem is also called Blind Source Separation (BSS). Though identifiability is impossible in the nonlinear case without further assumptions (Hyvärinen & Pajunen, 1999; Locatello et al., 2019), recent nonlinear ICA methods demonstrated identifiability results for increasingly realistic settings, pushing forward our understanding of when, how, and why SSL could work (Khemakhem et al., 2020b,a; Reizinger et al., 2022; Gresele et al., 2021; Morioka et al., 2021; Morioka & Hyvärinen, 2023; Hyvärinen & Morioka, 2016; Hyvärinen et al., 2019; Klindt et al., 2021). Identifiability guarantees are also central in the field of causal inference (Pearl, 2009). Causal methods are promising data efficiency and generalization, thus, they are of interest for practice. The new field of Causal Representation Learning (CRL) (Schölkopf et al., 2021) also has identifiability guarantees (von Kügelgen et al., 2023; von Kügelgen, 2024; Wendong et al., 2023; Brehmer et al., 2022; Ahuja et al., 2022), some of which were shown to be deeply connected to ICA methods (Reizinger et al., 2024b; Hyvärinen et al., 2023). Recently, Reizinger et al. (2024a) showed that even for standard classification tasks, a simple DGP can be formulated that enjoys linear identifiability. Moreover, this is easily estimated by maximum likelihood cross-entropy minimization, generalizing the result of Roeder et al. (2020) and providing evidence about why Huh et al. (2024) observed representational similarity across many training objectives—namely, they were different formulations of cross-entropy minimization. Despite delivering insights about widely-used SSL methods (Zimmermann et al., 2021), most theoretical works are fairly idealistic as they make strong assumptions, including assuming infinite data and infinite training time (i.e., model convergence).

### 3.4. Applications: why practitioners should care

The biggest critique of IT—and machine learning theories in general—is that 1) most progress in the field comes from empirical techniques or scaling data and compute and 2) theoretical results cannot realistically be transferred to practical scenarios (Sutton, 2019). We discuss this view in detail in § 5, but emphasize that algorithms with identifiable DGPs are used in many scientific and applied domains, including: robotics (Locatello et al., 2020; Lippe et al., 2023), dynamical systems (Lippe et al., 2022; Rajendran et al., 2023), neuroimaging (Himberg et al., 2004; Hyvärinen & Morioka,



2016), neuroscience (Zhou & Wei, 2020; Schneider et al., 2023), genomics (Morioka & Hyvarinen, 2023; Morioka & Hyvärinen, 2023), structural biology (Klindt et al., 2024), ant colonies (Dingling et al., 2024), and climate science (Yao et al., 2024). What is common in IT-driven applications is the focus on inferring the *correct* (unveiling the underlying science) and *robust* (generalizing, even OOD) mechanism from the data. Indeed, recent works in IT showed that OOD generalization is possible (Brady et al., 2023; 2025; Wiedemer et al., 2023a;b). Thus, these practically relevant considerations illustrate the benefit of IT-driven SSL design.

#### 4. Position: Singular Identifiability Theory as a blueprint to close the gap between SSL theory and practice

With the example of SimCLR, we showed how IT can improve our understanding of SSL methods (§ 3.2), particularly by making the assumptions on the data explicit via a DGP. However, IT relies on simplifications and cannot explain many practical phenomena, including dimensional collapse (Jing et al., 2022), or OOD behavior. *Thus, we need to move beyond the current paradigm and address the practical realities of modern machine learning. We call for an extension of IT to what we term Singular Identifiability Theory (SITH), and detail what it can bring to SSL.* Similarly to how Singular Learning Theory (SLT) (Watanabe, 2009; 2020) was proposed to address, among others, the reality that due to certain symmetries, (the Fisher information matrix of) neural networks can have singularities. We want to emphasize that SITH, in its current form, is an umbrella term for a *future* theory and a *blueprint* (Bronstein et al., 2021) for SSL research. We aim to provide signposts along which the path towards this new theory can be traversed. Thus, we review the gaps between SSL theory and practice (Tab. 1). We detail the empirical and/or theoretical evidence in the following subsections, pinpoint potential synergies, and formulate research questions, answering which, we believe, will move the field forward.

##### 4.1. The data augmentation gap: unrealistic DGPs

Extensive evaluations demonstrate that data augmentation strategy matters more than the SSL method (Morningstar et al., 2024). As we showed with SimCLR in § 3, the DGP is a useful construct to model the augmentations—and is at the core of identifiability results in Contrastive Learning (CL) (Hyvarinen & Morioka, 2016; Morioka et al., 2021; Zimmermann et al., 2021; Rusak et al., 2024). This means that changing an augmentation should be reflected in the DGP, and such changes might make some latents less feasible or impossible to learn (von Kügelgen et al., 2021; Daunhawer et al., 2023). The problem with current IT is that the modeling assumptions in the DGP are *too simplistic*. Using a hypersphere as latent space and simple

conditionals such as vMF distributions are neither principled nor realistic—for example, they cannot capture common augmentations such as (heavy) crops (Zimmermann et al., 2021; Rusak et al., 2024; Reizinger et al., 2024a).

**Question 4.1.1.** Moutakanni et al. (2024) suggest that prior knowledge is not required for augmentation design. To what extent is this the case, does this impose performance limitations, and is it possible to develop a theoretical understanding of this phenomenon?

##### 4.2. The asymptotic data gap: data set and batch size

Identifiability results assume infinite data, batch size, and converged, i.e., IT is an *asymptotic* theory—apart from Campi & Weyer (2002); Lyu & Fu (2022). In practice, data set size and data or task diversity (Elmoznino et al., 2024; Raventós et al., 2023) matters<sup>1</sup>. However, it is unclear whether data set size improves performance or it only correlates with properties like diversity in the sense of the ICA literature (intuitively, data coming from diverse environments), which is crucial for identifiability (Hyvarinen & Morioka, 2016; Morioka et al., 2021; Rajendran et al., 2023; Khemakhem et al., 2020a; Reizinger et al., 2024a). More data can *seemingly* improve OOD performance. However, Mayilvahanan et al. (2024a;b) showed that the real reason is that the OOD tasks became in-distribution as they were included in the training data. Batch size is also important in practice (Chen et al., 2020; Rusak et al., 2024), but there are no corresponding theoretical results. Recently, Reizinger et al. (2024a) showed that there is a connection between InfoNCE and DIET, where one difference is calculating the loss over the whole dataset for DIET, whereas InfoNCE uses a mini-batch of data. As both methods are identifiable, this might suggest that the role of batch size can be understood theoretically.

**Question 4.2.1.** Lyu & Fu (2022) provided a finite-sample analysis on the data diversity condition in nonlinear ICA, whereas Lyu et al. (2021) analyzed the sample complexity of recovering shared (content) latents in multi-view SSL. Can we use such results to provide practically meaningful recommendations on data set size and data diversity to achieve identifiable representations empirically?

**Question 4.2.2.** Is the number of tasks a model is trained on (called “task diversity” in Raventós et al. (2023)) a causal factor in improving model performance or is it more of a proxy to diversity in the ICA sense (Hyvarinen et al., 2019), i.e., more tasks presumably mean more diverse tasks? For example, Bansal et al. (2024) provides evidence that data diversity matters.

---

<sup>1</sup>Though “diversity” in these contexts often refers to the number of tasks, e.g., in Raventós et al. (2023)

Problem	Domain Status Quo	RQ's
<b>Data augmentations</b>	Theory ? Augmentations are modeled in the DGP, not in observation space Practice ✓ The role of augmentations is understood	4.1.1
<b>Finite data</b>	Theory ✗ Finite-sample analysis is almost entirely missing from IT Practice ✓ Scaling laws show the effect of data set size	4.2.1
<b>Data diversity</b>	Theory ✓ Data diversity is well-defined in the DGP view of IT Practice ✗ Data diversity is sometimes entangled with data set size	4.2.2
<b>Finite time</b>	Theory ? Crude understanding of (linear) training dynamics, but not tied to IT Practice ✓ Convergence differs across models	4.3.1 4.3.2
<b>Loss saturation</b>	Theory ? IT does not consider time, learning dynamics provides insights for linear models Practice ✗ It is unclear why convergence speeds differ	4.3.2
<b>Inductive biases</b>	Theory ✗ Results missing or not reflecting practical choices Practice ✓ The role of architecture is well-understood	4.4.1
<b>Initialization</b>	Theory ✗ IT does not consider initialization Practice ✓ It is well understood what initialization are useful	4.4.1
<b>Dimensional collapse</b>	Theory ✓ The DGP view defines what latents collapse Practice ? Computational tricks and regularizers might help collapse	4.5.1
<b>Projector</b>	Theory ✗ No explicit result, some settings reflect a linear projector Practice ? The projector is not eliminated, at most ameliorated	4.5.1
<b>Compositionality</b>	Theory ✓ Theoretical results exist for compositional generalization Practice ✗ Empirical methods struggle to generalize compositionally	4.6.1
<b>CL/non-CL</b>	Theory ✗ Theory for non-CL methods is missing Practice ✓ Representations are similar	4.7.1 4.7.2
<b>Evaluation</b>	Theory ✓ IT characterizes what latents are learned Practice ✗ Mostly benchmark-related, principles often missing	4.8.1 4.8.2

Table 1. The gaps between SSL theory and practice: RQ is a shorthand for research question, ✓ denotes a comprehensive understanding; ? incomplete or related results; and ✗ mostly uncharted territory.

### 4.3. The finite time gap: learning dynamics and loss saturation

Identifiability theory cannot distinguish between the convergence speed of models, as it focuses on models at the global optimum of the loss, i.e., at convergence. As experiments demonstrate, not achieving convergence to the global optimum—e.g., by a saturating loss, when it is very close to the optimum, but small differences can lead to qualitatively different models (Liu et al., 2023; Reizinger et al., 2024c)—is often the barrier to high-quality representations (Simon et al., 2023; Rusak et al., 2024; von Kügelgen et al., 2021; Kadkhodaie et al., 2024; Roeder et al., 2020). Another example of this is *grokking*, where the training loss saturates long before the true, generalizing solution is discovered (Power et al., 2022). Recently, theoretical insights started to emerge that provide some guidance on how to improve latent recovery via augmentation design, ensemble models, or hard negative sampling (Eastwood et al., 2023; Rusak et al.,

2024), which might guide further research understanding how such choices affect learning dynamics (potentially in terms of a DGP). Current IT aims to recover all latent factors, ignoring the minimality argument of Achille & Soatto (2018), and the practical considerations of efficiency and compression. “Partial” identifiability results such as those of von Kügelgen et al. (2021); Daunhawer et al. (2023) are generally seen in negative light—even though some latent factors might not matter downstream.

**Question 4.3.1.** Why is it the case that sometimes, even though identifiability is only provable for some latent factors, with a large enough latent space and training time, the other latent factors can also be learned (to some extent) (von Kügelgen et al., 2021)?

**Question 4.3.2.** Can we understand what model components, such as vanishing initialization (Simon et al., 2023; Kunin et al., 2024; Draganov et al., 2025), determine the convergence speed differences of SSL methods and how

data augmentation design and negative sampling change learning dynamics?

#### 4.4. The architecture gap: teacher/student networks, initialization, stop gradient

The SSL model zoo hinges on particular architectural choices like stop gradients and predictor networks, mostly in the non-contrastive domain (Bardes et al., 2021; Zbontar et al., 2021; Oquab et al., 2024). The necessity of engineering practices and model components needs to be critically evaluated to improve robustness, decrease computational cost, and to make theoretical analysis easier—leading to a potentially wider range of insights. The DIET (Ibrahim et al., 2024) instance discrimination method was proposed as a simplified SSL pipeline to avoid the theoretically not well-understood components of SSL (e.g., stop gradients). DIET performs comparably to other SSL methods and is proven to identify a cluster-centric DGP (Juhos et al., 2024; Reizinger et al., 2024a). These works demonstrated that SSL can provably work without many computational tricks, hopefully helping practitioners to design more efficient pipelines in the future.

**Question 4.4.1.** Morningstar et al. (2024) showed that many different algorithmic knobs do not substantially alter the downstream performance of self-supervised representations. How can we use this observation to develop theoretical results showing what are the necessary and sufficient architectural and data components to prove identifiability?

#### 4.5. The projector gap: dimensional collapse

The projector phenomenon refers to better latent recovery/downstream performance at the other-than-the-ultimate layers of the trained network (Chen et al., 2020). That is, some dimensions are collapsed (Jing et al., 2022). Note that this is distinct from *neural collapse* (Papayan et al., 2020) where the model learns one-hot predictions of the training data rather than learning the correct conditional probabilities as assumed, e.g., in Reizinger et al. (2024a). The projector is seen as “wasting” resources by training the last (few) layers, and its mitigation is actively researched (Tian, 2022; Bordes et al., 2023; Jing et al., 2022; Song et al., 2023; Bizeul et al., 2024; Xue et al., 2024; Saunshi et al., 2021). The projector is an accidental phenomenon—i.e., no one designs an SSL pipeline by saying “the last  $k$  layers are the projector, i.e., I expect them to be useless for downstream performance.” What is a deliberate choice is to use a linear or a nonlinear projector. Interestingly, a nonlinear projector is not necessarily better—see comparison in, e.g., (Mialon et al., 2022, Fig. 2). We argue that there is a possible theoretical explanation of both 1) why the projector phenomenon occurs, and 2) why a linear projector is often better than a nonlinear one. Roeder et al. (2020); Hyvarinen & Morioka (2016); Reizinger et al. (2024a) assumed that the trained neural network is a composition of a nonlinear encoder  $f$

and a linear projection matrix  $W$ , yielding  $W \circ f$ . Then they prove identifiability of  $z = f(x)$  in a log-linear model, where the  $z$  is fed through a softmax to yield the loss. We hypothesize that, in certain cases, this log-linear model can describe and explain the (linear) projector’s occurrence. So far, this potential link between identifiability results and the projector phenomenon has eluded the community, and requires further research.

**Question 4.5.1.** Roeder et al. (2020); Reizinger et al. (2024a) show that under a specific DGP, the key construct for identifiability is a log-linear model, i.e., estimating cross entropy with a softmax. Does this provably explain the (linear) projector phenomenon? Is it possible to develop an identifiability result without the projector?


#### 4.6. The generalization gap: OOD and compositionality

Identifiability results have recently been extended to specific OOD scenarios in computer vision, based on compositionality (Brady et al., 2023; 2025; Wiedemer et al., 2023b;a; Lachapelle et al., 2023). These results exploit that composition of attributes and objects is often straightforward to model in a DGP. Their proofs rely on imposing structural constraints such as additivity on the generative model. Most empirical representation learning methods do not generalize compositionally (Montero et al., 2021; 2024; Schott et al., 2021; Wiedemer et al., 2023a), emphasizing the need for theoretical insights to design better SSL methods. Interestingly, in the language domain, both the first theoretical results (Ahuja & Mansouri, 2024; Xie et al., 2022), and empirical techniques and observations (Bansal et al., 2022; Ruoss et al., 2023; Mészáros et al., 2024) exist showing that OOD generalization might be possible, at least in terms of length generalization or rule extrapolation.

**Question 4.6.1.** Is it possible to provide OOD or compositional generalization results across data modalities, such as in image-text models? If yes, is there a qualitative or quantitative difference w.r.t. the inductive bias of each modality?

#### 4.7. The Platonic representation gap: the contrastive and non-contrastive dichotomy

A prevalent categorization of SSL is the contrastive/non-contrastive dichotomy. While categorizations like these are undoubtedly valuable, they might obscure the overall trends within the SSL model landscape. However, recent initiatives are striving to bring unity to the field of SSL (Balestriero et al., 2023; Morningstar et al., 2024; Cabannes et al., 2023). Garrido et al. (2022) showed that both method families contrast some properties; for this reason, they propose the names *sample-contrastive* (for CL) and *dimension-contrastive* (for non-contrastive). Balestriero & LeCun (2022) analyzed contrastive and non-contrastive paradigms with a spectral embedding approach, showing that such methods correspond to specific local and global spectral embedding methods (e.g., SimCLR was shown to



**Figure 1. Representational similarities between contrastive and non-contrastive methods trained on ImageNet-1k:** (Left): the similarities are calculated with an RBF-kernel-based CKA with  $\sigma = 0.4$ . We include supervised baselines for reference. Reproduced from (Ciernik et al., 2024) with the authors’ permission. (Right): a two-dimensional UMAP embedding of similarity matrix colored by method types (contrastive, non-contrastive and supervised).

be equivalent with ISOMAP (Tenenbaum et al., 2000) under specific assumptions). Similarly, t-SNE can be made contrastive with minimal modifications (Böhm et al., 2022). There is also evidence that both paradigms relate to entropy and mutual information estimation (Liu et al., 2022). Wang & Isola (2020) interpreted InfoNCE as an interaction between uniformity (i.e., entropy estimation), and alignment (i.e., invariance w.r.t. augmentations). The InfoNCE family is shown to minimize cross-entropy (Zimmermann et al., 2021; Rusak et al., 2024); the same holds for the (non-contrastive) instance-discrimination method called DIET, which is also related to InfoNCE (Ibrahim et al., 2024; Reizinger et al., 2024a). These principles are also present in non-contrastive methods, though usually via regularizers. For example, VICReg (Bardes et al., 2021) also estimates a lower bound on entropy (Shwartz-Ziv et al., 2022). One distinction is the absence of an explicit DGP in most non-CL methods. That is, the DGP is mostly defined implicitly, with the exception of (Reizinger et al., 2024a).

The contrastive and non-contrastive split posits a sterile dichotomy

Contrastive and non-contrastive methods are driven by the same principles, thus, we need to consider their similarities, and not focus on their differences as they represent different means to the same end.

Experimental evidence demonstrates that (identifiable) contrastive and non-contrastive representations fare similarly on some downstream task (Morningstar et al., 2024) and that their representations can be highly similar (Ciernik et al., 2024)—cf. Fig. 1 and also note similarity is not universal, e.g., ViT-based models are less similar. However, identifiability guarantees for non-contrastive methods are mostly

missing—with the exception of DIET (Ibrahim et al., 2024; Reizinger et al., 2024a). The empirical similarity suggests that some non-contrastive methods might be identifiable.

**Question 4.7.1.** Based on the similarity of contrastive and non-contrastive representations (Morningstar et al., 2024; Ciernik et al., 2024) and the identifiability of contrastive methods, can we prove identifiability for non-contrastive methods?

**Question 4.7.2.** The representations of contrastive and non-contrastive methods are often similar (Ciernik et al., 2024) (cf. Fig. 1), though their convergence speed is highly distinct (Simon et al., 2023). Can we theoretically understand in which (data, model size, etc.) regimes which method is expected to perform best?

#### 4.8. The evaluation gap: going beyond ImageNet

Most SSL methods evaluate performance on ImageNet (Krizhevsky et al., 2012) downstream classification, without acknowledging that such an evaluation can only falsify claims about classification-related latents. For example, if orientation is not required for classification, then a more universal representation (e.g., one capturing orientation) will not perform better—though this makes a difference in terms of “universality” of the representation. Indeed, SSL methods capturing more latents sometimes have (slightly) lower downstream classification accuracy (Rusak et al., 2024), which might suggest their infeasibility—though that depends on what they will be used for. *We need to question the practice of measuring the “universality” of representation with a single scalar on a highly specialized task.* A good example of a more comprehensive evaluation is DINOv2 (Oquab et al., 2024). Another gap, especially for evaluating identifiability claims, is the lack of large-scale data sets with latent information—proxies such as



ImageNet-X (Idrissi et al., 2022) exist, though they only have binary labels. Recent works started using aggregate statistics (e.g., rank conditions) to evaluate the representation quality (Agrawal et al., 2022; Garrido et al., 2023; Thilak et al., 2023). These methods seem to be somewhat predictive of even OOD performance, we lack a theoretical understanding when and why this is the case.

Evaluating self-supervised representations requires more than ImageNet classification

SITh can improve SSL evaluations by using a DGP to focus on the relevant latent factors and design more principled benchmarks.

**Question 4.8.1.** How can we evaluate the “universality” of self-supervised representations for robustness? I.e., how should we move away from downstream classification on ImageNet? What are the *set* of downstream tasks we need to evaluate SSL methods?

**Question 4.8.2.** How can we develop ImageNet-scale (synthetic) datasets with latent information to evaluate the identifiability of self-supervised representations? Can we use rendering pipelines akin to DisLib (Locatello et al., 2019) to develop more principled benchmarks?

#### 4.9. Position summary: We need an empirically grounded identifiability theory for SSL

To improve our understanding of the (implicit) operating principles of SSL methods such as SimCLR (Zimmermann et al., 2021) (§ 3.2). Recent IT results also improved our understanding of compositional generalization (Brady et al., 2023; 2025; Wiedemer et al., 2023b;a; Lachapelle et al., 2023). Yet, current IT cannot fully explain empirical observations in SSL. To quote Nobel-prize-winning physicist Richard P. Feynman: *“It doesn’t matter how beautiful your theory is, if it doesn’t agree with the experiment, it’s wrong.”* IT models the data via an underlying data generating process (DGP), however, the design of the DGP does not exploit guidance from empirical observations, which leads to presumably the biggest critique of identifiability theory: *Identifiability guarantees are not valuable per se, only to the extent they can describe empirical phenomena, or—and only time will tell this—they aid the development of such theories that explain such phenomena.*

When a theorist constructs a DGP, the focus should not only be on identifiability but also on the match with reality (Klindt et al., 2021). Vice versa, when a practitioner sets up a pre-text task for SSL, they should think about the implicit DGP that they might be assuming and turn to theory to ask whether this setup results in an identifiable DGP (and estimation procedure). Constructing a DGP on the hypersphere is common in theory (Zimmermann et al., 2021; Reizinger et al., 2024a). However, it is unclear how well

this actually corresponds to a realistic DGP (for images). Analogously, building a pre-text around data augmentations is an empirically successful construct, however, it will only result in theoretic identifiability if the data augmentations cover all degrees of freedom in the data—a rather unrealistic assumption, mirrored by findings on the importance of the right augmentations (Chen et al., 2020; Oquab et al., 2024; Morningstar et al., 2024).

## 5. Alternative Views

Recent years’ breakthroughs in AI were often fueled by scaling up models, compute, and data (Sutton, 2019). Though empirical algorithmic improvements are also crucial for breakthroughs like the recent DeepSeek-R1 (DeepSeek-AI et al., 2025) reasoning model. The scaling-related observations were documented in so-called scaling laws (Hernandez et al., 2021; Zhai et al., 2022; Henighan et al., 2020; Hestness et al., 2017; Kaplan et al., 2020), analyzing trends in loss reduction. Especially in the case of Large Language Models (LLMs), scaling data and compute is still significant for progress, maintaining the popularity of this opinion in the community in general.

**Counterarguments in SSL:** Realistically, engineering efforts are crucial to solve the problems of real-world SSL systems. However, we are seeing the SSL community shift beyond scaling.<sup>2</sup> After scraping the majority of data from the Internet, and pushing the boundaries of compute beyond previously unimaginable limits, we need a new approach. Sorscher et al. (2022) theoretically showed that scaling laws can sometimes be overcome via pruning. More data can also imply that OOD tasks become in-distribution (Mayilvahanan et al., 2024a;b), qualitatively changing what we as a community thought about the true role of scaling data. We propose an identifiability theoretical framework as an alternative to the scaling view, still prevalent in LLM research. We believe that having a more principled approach towards SSL, which incorporates the empirical observations, is the way to move SSL forward. Particularly, improving OOD generalization performance in SSL is hard to do empirically. First, most empirical representation learning methods, even despite aiming to learn disentangled representations, do not generalize compositionally (Montero et al., 2021; 2024; Schott et al., 2021; Wiedemer et al., 2023a). Second, it is hard to get the intuition without using a model like a DGP. We believe that Singular Identifiability Theory (SITh) can help overcome these limitations, and lead to more efficient, principled, and robust SSL methods.

## 6. Conclusion

Our position paper hit a critical note to highlight what we conceive as a big barrier before closing the gap between

<sup>2</sup>Personal observations at the NeurIPS 2024 Workshop: Self-Supervised Learning - Theory and Practice



SSL theory and practice. We argued that thinking in terms of an empirically-grounded data generating process (DGP) can accelerate SSL research. We formulated our suggestion in terms of an extension to identifiability theory, called Singular Identifiability Theory (SITh), to foster synergies between the two subfields. To provide a common starting point, we included an intuitive introduction of identifiability for practitioners (§ 3), demonstrating how SSL practice can benefit from IT. However, more research needs to be done, illustrated by the many open questions in both theory and practice (Tab. 1). As we are close to hitting the limits of scaling-driven advances, we need to turn towards more principled, theoretically grounded approaches. This is especially important when machine learning models are deployed in safety-critical environments such as autonomous driving or healthcare. We believe that SITh facilitates asking better research questions, many of which we shared, to drive SSL forward by improving how we design, evaluate, and understand SSL algorithms.

## Acknowledgements

The authors would like to thank Laure Ciernik, Evgenia Rusak, Roland Zimmermann, Julius von Kügelgen, Michael Kirchhof, Lucas Maes, Vishaal Udandarao, Ameya Prabhu, Karsten Roth, Omar Chehab, Runtian Zhai, Kamilė Stankevičiūtė, Thaddäus Wiedemer, Pradeep Ravikumar, Luigi Gresle, Quentin Garrido, Alice Bizeul, Mark Ibrahim, and Attila Juhos for their valuable suggestions for shaping the manuscript. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Patrik Reizinger. Patrik Reizinger acknowledges his membership in the European Laboratory for Learning and Intelligent Systems (ELLIS) PhD program. This work was supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A. Wieland Brendel acknowledges financial support via an Emmy Noether Grant funded by the German Research Foundation (DFG) under grant no. BR 6382/1-1 and via the Open Philanthropy Foundation funded by the Good Ventures Foundation. Wieland Brendel is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645. This research utilized compute resources at the Tübingen Machine Learning Cloud, DFG FKZ INST 37/1057-1 FUGG.

## Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

## References

- Achille, A. and Soatto, S. Emergence of Invariance and Disentanglement in Deep Representations. In *2018 Information Theory and Applications Workshop (ITA)*, pp. 1–9, San Diego, CA, February 2018. IEEE. ISBN 978-1-72810-124-8. doi: 10.1109/ITA.2018.8503149. URL <https://ieeexplore.ieee.org/document/8503149/>. 5
- Agrawal, K. K., Mondal, A. K., Ghosh, A., and Richards, B. A.  $\alpha$ -ReQ : Assessing  $\text{ReQ}$  representation quality in Self-Supervised Learning by measuring eigenspectrum decay. October 2022. URL <https://openreview.net/forum?id=ii9X4vtZGTZ>. 8
- Ahuja, K. and Mansouri, A. On Provable Length and Compositional Generalization, February 2024. URL <http://arxiv.org/abs/2402.04875>. arXiv:2402.04875 [cs, stat]. 6
- Ahuja, K., Wang, Y., Mahajan, D., and Bengio, Y. Interventional Causal Representation Learning, September 2022. URL <http://arxiv.org/abs/2209.11924>. arXiv:2209.11924 [cs, stat]. 3, 18
- Arora, S., Li, Y., Liang, Y., Ma, T., and Risteski, A. A Latent Variable Model Approach to PMI-based Word Embeddings. *Transactions of the Association for Computational Linguistics*, 4:385–399, July 2016. ISSN 2307-387X. doi: 10.1162/tacl.a.00106. URL <https://doi.org/10.1162/tacl.a.00106>. 2
- Assran, M., Duval, Q., Misra, I., Bojanowski, P., Vincent, P., Rabbat, M., LeCun, Y., and Ballas, N. Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture, April 2023. URL <http://arxiv.org/abs/2301.08243>. arXiv:2301.08243 [cs, eess]. 1
- Balestriero, R. and LeCun, Y. Contrastive and Non-Contrastive Self-Supervised Learning Recover Global and Local Spectral Embedding Methods, June 2022. URL <http://arxiv.org/abs/2205.11508>. arXiv:2205.11508 [cs, math, stat]. 6
- Balestriero, R., Ibrahim, M., Sobal, V., Morcos, A., Shekhar, S., Goldstein, T., Bordes, F., Bardes, A., Mialon, G., Tian, Y., Schwarzschild, A., Wilson, A. G., Geiping, J., Garrido, Q., Fernandez, P., Bar, A., Pirsiavash, H., LeCun, Y., and Goldblum, M. A Cookbook of Self-Supervised Learning, April 2023. URL <http://arxiv.org/abs/2304.12210>. arXiv:2304.12210 [cs]. 2, 6, 18
- Bansal, A., Schwarzschild, A., Borgnia, E., Emam, Z., Huang, F., Goldblum, M., and Goldstein, T. End-to-end Algorithm Synthesis with Recurrent Networks: Logical Extrapolation Without Overthinking, October 2022. URL <http://arxiv.org/abs/2202.05826>. arXiv:2202.05826 [cs]. 6

- Bansal, H., Hosseini, A., Agarwal, R., Tran, V. Q., and Kazemi, M. Smaller, Weaker, Yet Better: Training LLM Reasoners via Compute-Optimal Sampling, October 2024. URL <http://arxiv.org/abs/2408.16737>. arXiv:2408.16737 [cs]. 4
- Bardes, A., Ponce, J., and LeCun, Y. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning. *arXiv:2105.04906 [cs]*, May 2021. URL <http://arxiv.org/abs/2105.04906>. arXiv: 2105.04906. 6, 7
- Bizeul, A., Schölkopf, B., and Allen, C. A Probabilistic Model to explain Self-Supervised Representation Learning, February 2024. URL <http://arxiv.org/abs/2402.01399>. arXiv:2402.01399 [cs, stat]. 1, 6
- Bloem-Reddy, B. and Teh, Y. W. Probabilistic symmetries and invariant neural networks. *The Journal of Machine Learning Research*, 21(1):90:3535–90:3595, January 2020. ISSN 1532-4435. 18
- Böhm, J. N., Berens, P., and Kobak, D. Unsupervised visualization of image datasets using contrastive learning. *arXiv preprint arXiv:2210.09879*, 2022. 7
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosse-lut, A., Brunskill, E., et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 2
- Bordes, F., Balestriero, R., Garrido, Q., Bardes, A., and Vincent, P. Guillotine Regularization: Why removing layers is needed to improve generalization in Self-Supervised Learning. *Transactions on Machine Learning Research*, May 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=ZgXfXSz5ln&s=09>. 6
- Brady, J., Zimmermann, R. S., Sharma, Y., Schölkopf, B., von Kügelgen, J., and Brendel, W. Provably Learning Object-Centric Representations, May 2023. URL <http://arxiv.org/abs/2305.14229>. arXiv:2305.14229 [cs]. 4, 6, 8, 18
- Brady, J., von Kügelgen, J., Lachapelle, S., Buchholz, S., Kipf, T., and Brendel, W. Interaction asymmetry: A general principle for learning composable abstractions. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=cCl10IU836>. 4, 6, 8, 18
- Brehmer, J., de Haan, P., Lippe, P., and Cohen, T. Weakly supervised causal representation learning, October 2022. URL <http://arxiv.org/abs/2203.16437>. arXiv:2203.16437 [cs, stat]. 3
- Bronstein, M. M., Bruna, J., Cohen, T., and Velicković, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021. 1, 4
- Buchholz, S., Besserve, M., and Schölkopf, B. Function Classes for Identifiable Nonlinear Independent Component Analysis. pp. 25, 2022. URL <https://arxiv.org/abs/2208.06406>. 18
- Cabannes, V., Kiani, B. T., Balestriero, R., LeCun, Y., and Bietti, A. The SSL Interplay: Augmentations, Inductive Bias, and Generalization, February 2023. URL <http://arxiv.org/abs/2302.02774>. arXiv:2302.02774 [cs, math, stat]. 6, 18
- Campi, M. and Weyer, E. Finite sample properties of system identification methods. *IEEE Transactions on Automatic Control*, 47(8):1329–1334, 2002. doi: 10.1109/TAC.2002.800750. 4
- Cannistraci, I., Moschella, L., Fumero, M., Maiorca, V., and Rodolà, E. From Bricks to Bridges: Product of Invariances to Enhance Latent Space Communication, October 2023. URL <http://arxiv.org/abs/2310.01211>. arXiv:2310.01211 [cs]. 2
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A Simple Framework for Contrastive Learning of Visual Representations. *arXiv:2002.05709 [cs, stat]*, June 2020. URL <http://arxiv.org/abs/2002.05709>. arXiv: 2002.05709. 2, 3, 4, 6, 8, 18
- Ciernik, L., Linhardt, L., Morik, M., Dippel, J., Kornblith, S., and Muttenthaler, L. Training objective drives the consistency of representational similarity across datasets, November 2024. URL <http://arxiv.org/abs/2411.05561>. arXiv:2411.05561 [cs]. 2, 7, 18
- Comon, P. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994. 3
- Daunhawer, I., Bizeul, A., Palumbo, E., Marx, A., and Vogt, J. E. Identifiability Results for Multimodal Contrastive Learning, March 2023. URL <http://arxiv.org/abs/2303.09166>. arXiv:2303.09166 [cs, stat] version: 1. 4, 5
- DeepSeek-AI, Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., Zhang, X., Yu, X., Wu, Y., Wu, Z. F., Gou, Z., Shao, Z., Li, Z., Gao, Z., Liu, A., Xue, B., Wang, B., Wu, B., Feng, B., Lu, C., Zhao, C., Deng, C., Zhang, C., Ruan, C., Dai, D., Chen, D., Ji, D., Li, E., Lin, F., Dai, F., Luo, F., Hao, G., Chen, G., Li, G., Zhang, H., Bao, H., Xu, H., Wang, H., Ding, H., Xin, H., Gao, H., Qu, H., Li, H., Guo, J., Li, J., Wang, J., Chen, J., Yuan, J., Qiu, J.,

- Li, J., Cai, J. L., Ni, J., Liang, J., Chen, J., Dong, K., Hu, K., Gao, K., Guan, K., Huang, K., Yu, K., Wang, L., Zhang, L., Zhao, L., Wang, L., Zhang, L., Xu, L., Xia, L., Zhang, M., Zhang, M., Tang, M., Li, M., Wang, M., Li, M., Tian, N., Huang, P., Zhang, P., Wang, Q., Chen, Q., Du, Q., Ge, R., Zhang, R., Pan, R., Wang, R., Chen, R. J., Jin, R. L., Chen, R., Lu, S., Zhou, S., Chen, S., Ye, S., Wang, S., Yu, S., Zhou, S., Pan, S., Li, S. S., Zhou, S., Wu, S., Ye, S., Yun, T., Pei, T., Sun, T., Wang, T., Zeng, W., Zhao, W., Liu, W., Liang, W., Gao, W., Yu, W., Zhang, W., Xiao, W. L., An, W., Liu, X., Wang, X., Chen, X., Nie, X., Cheng, X., Liu, X., Xie, X., Liu, X., Yang, X., Li, X., Su, X., Lin, X., Li, X. Q., Jin, X., Shen, X., Chen, X., Sun, X., Wang, X., Song, X., Zhou, X., Wang, X., Shan, X., Li, Y. K., Wang, Y. Q., Wei, Y. X., Zhang, Y., Xu, Y., Li, Y., Zhao, Y., Sun, Y., Wang, Y., Yu, Y., Zhang, Y., Shi, Y., Xiong, Y., He, Y., Piao, Y., Wang, Y., Tan, Y., Ma, Y., Liu, Y., Guo, Y., Ou, Y., Wang, Y., Gong, Y., Zou, Y., He, Y., Xiong, Y., Luo, Y., You, Y., Liu, Y., Zhou, Y., Zhu, Y. X., Xu, Y., Huang, Y., Li, Y., Zheng, Y., Zhu, Y., Ma, Y., Tang, Y., Zha, Y., Yan, Y., Ren, Z. Z., Ren, Z., Sha, Z., Fu, Z., Xu, Z., Xie, Z., Zhang, Z., Hao, Z., Ma, Z., Yan, Z., Wu, Z., Gu, Z., Zhu, Z., Liu, Z., Li, Z., Xie, Z., Song, Z., Pan, Z., Huang, Z., Xu, Z., Zhang, Z., and Zhang, Z. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL <https://arxiv.org/abs/2501.12948>. 8
- Dingling, Y., Rancati, D., Cadei, R., Fumero, M., and Locatello, F. Unifying causal representation learning with the invariance principle. *arXiv preprint arXiv:2409.02772*, 2024. 4
- Draganov, A., Vadgama, S., Damrich, S., Böhm, J. N., Maes, L., Kobak, D., and Bekkers, E. On the Importance of Embedding Norms in Self-Supervised Learning, February 2025. URL <http://arxiv.org/abs/2502.09252>. arXiv:2502.09252 [cs]. 5
- Duan, S., Matthey, L., Saraiva, A., Watters, N., Burgess, C. P., Lerchner, A., and Higgins, I. Unsupervised Model Selection for Variational Disentangled Representation Learning. *arXiv:1905.12614 [cs, stat]*, February 2020. URL <http://arxiv.org/abs/1905.12614>. arXiv: 1905.12614. 2
- Eastwood, C., von Kügelgen, J., Ericsson, L., Bouchacourt, D., Vincent, P., Schölkopf, B., and Ibrahim, M. Self-Supervised Disentanglement by Leveraging Structure in Data Augmentations, November 2023. URL <http://arxiv.org/abs/2311.08815>. arXiv:2311.08815 [cs, stat]. 5
- Elmoznino, E., Marty, T., Kasetty, T., Gagnon, L., Mittal, S., Fathi, M., Sridhar, D., and Lajoie, G. In-context learning and Occam’s razor, October 2024. URL <http://arxiv.org/abs/2410.14086>. arXiv:2410.14086. 4
- Fleissner, M., Esser, P., and Ghoshdastidar, D. A Probabilistic Model for Self-Supervised Learning, January 2025. URL <http://arxiv.org/abs/2501.13031>. arXiv:2501.13031 [cs]. 1
- Fumero, M., Pegoraro, M., Maiorca, V., Locatello, F., and Rodolà, E. Latent Functional Maps: a spectral framework for representation alignment. 2024. 2
- Garrido, Q., Chen, Y., Bardes, A., Najman, L., and Lecun, Y. On the duality between contrastive and non-contrastive self-supervised learning, October 2022. URL <http://arxiv.org/abs/2206.02574>. arXiv:2206.02574 [cs]. 6
- Garrido, Q., Balestrieri, R., Najman, L., and Lecun, Y. Rankme: Assessing the downstream performance of pretrained self-supervised representations by their rank, 2023. URL <https://arxiv.org/abs/2210.02885>. 8
- Girdhar, R., Singh, M., Ravi, N., Van Der Maaten, L., Joulin, A., and Misra, I. Omnivore: A single model for many visual modalities. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16102–16112, 2022. 2
- Gresele, L., von Kügelgen, J., Stimper, V., Schölkopf, B., and Besserve, M. Independent mechanism analysis, a new concept? *arXiv:2106.05200 [cs, stat]*, June 2021. URL <http://arxiv.org/abs/2106.05200>. arXiv: 2106.05200. 3, 18
- Henighan, T., Kaplan, J., Katz, M., Chen, M., Hesse, C., Jackson, J., Jun, H., Brown, T. B., Dhariwal, P., Gray, S., Hallacy, C., Mann, B., Radford, A., Ramesh, A., Ryder, N., Ziegler, D. M., Schulman, J., Amodei, D., and McCandlish, S. Scaling Laws for Autoregressive Generative Modeling, November 2020. URL <http://arxiv.org/abs/2010.14701>. arXiv:2010.14701 [cs]. 8
- Hernandez, D., Kaplan, J., Henighan, T., and McCandlish, S. Scaling Laws for Transfer, February 2021. URL <http://arxiv.org/abs/2102.01293>. arXiv:2102.01293 [cs]. 8
- Hestness, J., Narang, S., Ardalani, N., Diamos, G., Jun, H., Kianinejad, H., Patwary, M. M. A., Yang, Y., and Zhou, Y. Deep Learning Scaling is Predictable, Empirically, December 2017. URL <http://arxiv.org/abs/1712.00409>. arXiv:1712.00409 [cs]. 8

- Higgins, I., Amos, D., Pfau, D., Racaniere, S., Matthey, L., Rezende, D., and Lerchner, A. Towards a Definition of Disentangled Representations. *arXiv:1812.02230 [cs, stat]*, December 2018. URL <http://arxiv.org/abs/1812.02230>. arXiv: 1812.02230. 3
- Himberg, J., Hyvärinen, A., and Esposito, F. Validating the independent components of neuroimaging time series via clustering and visualization. *Neuroimage*, 22(3):1214–1222, 2004. 3
- Hoyer, P., Janzing, D., Mooij, J. M., Peters, J., and Schölkopf, B. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2008. URL <https://proceedings.neurips.cc/paper/2008/hash/f7664060cc52bc6f3d620bcedc94a4b6-Abstract.html>. 18
- Huh, M., Cheung, B., Wang, T., and Isola, P. The Platonic Representation Hypothesis, May 2024. URL <http://arxiv.org/abs/2405.07987>. arXiv:2405.07987 [cs]. 1, 2, 3, 18
- Hyvarinen, A. and Morioka, H. Unsupervised Feature Extraction by Time-Contrastive Learning and Nonlinear ICA. *arXiv:1605.06336 [cs, stat]*, May 2016. URL <http://arxiv.org/abs/1605.06336>. arXiv: 1605.06336. 3, 4, 6, 18
- Hyvarinen, A. and Morioka, H. Nonlinear ICA of Temporally Dependent Stationary Sources. In *Artificial Intelligence and Statistics*, pp. 460–469. PMLR, April 2017. URL <http://proceedings.mlr.press/v54/hyvarinen17a.html>. ISSN: 2640-3498. 3
- Hyvarinen, A., Karhunen, J., and Oja, E. *Independent component analysis*. J. Wiley, New York, 2001. ISBN 978-0-471-40540-5. 3
- Hyvarinen, A., Sasaki, H., and Turner, R. E. Nonlinear ICA Using Auxiliary Variables and Generalized Contrastive Learning. *arXiv:1805.08651 [cs, stat]*, February 2019. URL <http://arxiv.org/abs/1805.08651>. arXiv: 1805.08651. 3, 4, 18
- Hyvärinen, A. Independent component analysis: recent advances. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1984):20110534, February 2013. ISSN 1364-503X, 1471-2962. doi: 10.1098/rsta.2011.0534. URL <https://royalsocietypublishing.org/doi/10.1098/rsta.2011.0534>. 18
- Hyvärinen, A. and Pajunen, P. Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 12(3):429–439, April 1999. ISSN 0893-6080. doi: 10.1016/S0893-6080(98)00140-3. URL <https://www.sciencedirect.com/science/article/pii/S0893608098001403>. 3, 18
- Hyvärinen, A., Khemakhem, I., and Monti, R. Identifiability of latent-variable and structural-equation models: from linear to nonlinear, February 2023. URL <http://arxiv.org/abs/2302.02672>. arXiv:2302.02672 [cs, stat]. 3, 18
- Ibrahim, M., Klindt, D., and Balestrieri, R. Occam’s Razor for Self Supervised Learning: What is Sufficient to Learn Good Representations?, June 2024. URL <http://arxiv.org/abs/2406.10743>. arXiv:2406.10743 [cs]. 6, 7
- Idrissi, B. Y., Bouchacourt, D., Balestrieri, R., Evtimov, I., Hazirbas, C., Ballas, N., Vincent, P., Drozdal, M., Lopez-Paz, D., and Ibrahim, M. Imagenet-x: Understanding model mistakes with factor of variation annotations, 2022. URL <https://arxiv.org/abs/2211.01866>. 8
- Jha, R., Zhang, C., Shmatikov, V., and Morris, J. X. Harnessing the Universal Geometry of Embeddings, May 2025. URL <http://arxiv.org/abs/2505.12540>. arXiv:2505.12540 [cs]. 2
- Jing, L., Vincent, P., LeCun, Y., and Tian, Y. Understanding Dimensional Collapse in Contrastive Self-supervised Learning, April 2022. URL <http://arxiv.org/abs/2110.09348>. Number: arXiv:2110.09348 arXiv:2110.09348 [cs]. 1, 2, 3, 4, 6, 18
- Juhos, A., Bizeul, A., Reizinger, P., Balestrieri, R., Klindt, David, Ibrahim, M., Vogt, J. E., and Brendel, W. Dieting: Self-supervised learning with instance discrimination learns identifiable features. *NeurIPS 2024 Workshop on Self-Supervised Learning - Theory and Practice*, 2024. 6
- Kadkhodaie, Z., Guth, F., Simoncelli, E. P., and Mallat, S. Generalization in diffusion models arises from geometry-adaptive harmonic representations. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=ANvmVS2Yr0>. 5
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. Scaling Laws for Neural Language Models, January 2020. URL <http://arxiv.org/abs/2001.08361>. arXiv:2001.08361 [cs]. 8
- Khemakhem, I., Kingma, D., Monti, R., and Hyvarinen, A. Variational Autoencoders and Nonlinear ICA: A Unifying Framework. In *International Conference on Artificial*



- Intelligence and Statistics*, pp. 2207–2217. PMLR, June 2020a. URL <http://proceedings.mlr.press/v108/khemakhem20a.html>. ISSN: 2640-3498. 3, 4
- Khemakhem, I., Monti, R. P., Kingma, D. P., and Hyvärinen, A. ICE-BeeM: Identifiable Conditional Energy-Based Deep Models Based on Nonlinear ICA. *arXiv:2002.11537 [cs, stat]*, October 2020b. URL <http://arxiv.org/abs/2002.11537>. arXiv: 2002.11537. 3
- 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. *arXiv:2007.10930 [cs, stat]*, March 2021. URL <http://arxiv.org/abs/2007.10930>. arXiv: 2007.10930. 3, 8
- Klindt, D. A., Hyvärinen, A., Levy, A., Miolane, N., and Poitevin, F. Towards interpretable cryo-em: disentangling latent spaces of molecular conformations. *Frontiers in Molecular Biosciences*, 11:1393564, 2024. 4
- Krishnan, R., Rajpurkar, P., and Topol, E. J. Self-supervised learning in medicine and healthcare. *Nature Biomedical Engineering*, 6(12):1346–1352, 2022. 2
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 2, 7
- Kulkarni, T. D., Whitney, W. F., Kohli, P., and Tenenbaum, J. Deep convolutional inverse graphics network. *Advances in neural information processing systems*, 28, 2015. 2
- Kunin, D., Raventós, A., Dominé, C., Chen, F., Klindt, D., Saxe, A., and Ganguli, S. Get rich quick: exact solutions reveal how unbalanced initializations promote rapid feature learning, October 2024. URL <http://arxiv.org/abs/2406.06158>. arXiv:2406.06158. 5
- Lachapelle, S., Brouillard, P., Deleu, T., and Lacoste-Julien, S. Gradient-Based Neural DAG Learning. *arXiv:1906.02226 [cs, stat]*, February 2020. URL <http://arxiv.org/abs/1906.02226>. arXiv: 1906.02226. 18
- Lachapelle, S., Mahajan, D., Mitliagkas, I., and Lacoste-Julien, S. Additive Decoders for Latent Variables Identification and Cartesian-Product Extrapolation, July 2023. URL <http://arxiv.org/abs/2307.02598>. arXiv:2307.02598 [cs, stat]. 6, 8, 18
- Lippe, P., Magliacane, S., Löwe, S., Asano, Y. M., Cohen, T., and Gavves, E. iCITRIS: Causal Representation Learning for Instantaneous Temporal Effects. July 2022. URL <https://openreview.net/forum?id=xekDTZs7Z7>. 3
- Lippe, P., Magliacane, S., Löwe, S., Asano, Y. M., Cohen, T., and Gavves, E. BISCUIT: Causal Representation Learning from Binary Interactions, June 2023. URL <http://arxiv.org/abs/2306.09643>. arXiv:2306.09643 [cs, stat]. 3
- Liu, H., Xie, S. M., Li, Z., and Ma, T. Same Pre-training Loss, Better Downstream: Implicit Bias Matters for Language Models. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 22188–22214. PMLR, July 2023. URL <https://proceedings.mlr.press/v202/liu23ao.html>. ISSN: 2640-3498. 5
- Liu, X., Wang, Z., Li, Y., and Wang, S. Self-Supervised Learning via Maximum Entropy Coding. 2022. doi: 10.48550/ARXIV.2210.11464. URL <https://arxiv.org/abs/2210.11464>. Publisher: arXiv Version Number: 1. 7
- Locatello, F., Bauer, S., Lucic, M., Raetsch, G., Gelly, S., Schölkopf, B., and Bachem, O. Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations. In *International Conference on Machine Learning*, pp. 4114–4124. PMLR, May 2019. URL <http://proceedings.mlr.press/v97/locatello19a.html>. ISSN: 2640-3498. 3, 8, 18
- Locatello, F., Poole, B., Rätsch, G., Schölkopf, B., Bachem, O., and Tschannen, M. Weakly-Supervised Disentanglement Without Compromises. *arXiv:2002.02886 [cs, stat]*, October 2020. URL <http://arxiv.org/abs/2002.02886>. arXiv: 2002.02886. 3
- Lyu, Q. and Fu, X. On Finite-Sample Identifiability of Contrastive Learning-Based Nonlinear Independent Component Analysis. In *Proceedings of the 39th International Conference on Machine Learning*, pp. 14582–14600. PMLR, June 2022. URL <https://proceedings.mlr.press/v162/lyu22a.html>. ISSN: 2640-3498. 4
- Lyu, Q., Fu, X., Wang, W., and Lu, S. Understanding Latent Correlation-Based Multiview Learning and Self-Supervision: An Identifiability Perspective. October 2021. URL <https://openreview.net/forum?id=5FUq05QRc5b>. 4
- Mahajan, D., Pezeshki, M., Mitliagkas, I., Ahuja, K., and Vincent, P. Compositional Risk Minimization, October 2024. URL <http://arxiv.org/abs/2410.06303>. arXiv:2410.06303. 18
- Maiorca, V., Moschella, L., Fumero, M., Locatello, F., and Rodolà, E. Latent Space Translation via Inverse Relative Projection, June 2024. URL <http://arxiv.org/abs/2406.15057>. arXiv:2406.15057 [cs]. 2

- Marconato, E., Lachapelle, S., Weichwald, S., and Gresele, L. All or None: Identifiable Linear Properties of Next-token Predictors in Language Modeling, October 2024. URL <http://arxiv.org/abs/2410.23501>. arXiv:2410.23501. 2
- Mayilvahanan, P., Wiedemer, T., Rusak, E., Bethge, M., and Brendel, W. Does CLIP’s Generalization Performance Mainly Stem from High Train-Test Similarity?, March 2024a. URL <http://arxiv.org/abs/2310.09562>. arXiv:2310.09562 [cs]. 1, 4, 8
- Mayilvahanan, P., Zimmermann, R. S., Wiedemer, T., Rusak, E., Juhos, A., Bethge, M., and Brendel, W. In Search of Forgotten Domain Generalization, October 2024b. URL <http://arxiv.org/abs/2410.08258>. arXiv:2410.08258 [cs]. 1, 4, 8
- Mialon, G., Balestrieri, R., and LeCun, Y. Variance Covariance Regularization Enforces Pairwise Independence in Self-Supervised Representations, September 2022. URL <http://arxiv.org/abs/2209.14905>. arXiv:2209.14905 [cs]. 6
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL [https://papers.nips.cc/paper\\_files/paper/2013/hash/9aa42b31882ec039965f3c4923ce901b-Abstract.html](https://papers.nips.cc/paper_files/paper/2013/hash/9aa42b31882ec039965f3c4923ce901b-Abstract.html). 2
- Montero, M. L., Ludwig, C. J. H., Costa, R. P., Malhotra, G., and Bowers, J. S. THE ROLE OF DISENTANGLE-MENT IN GENERALISATION. pp. 17, 2021. 1, 6, 8
- Montero, M. L., Bowers, J. S., Costa, R. P., Ludwig, C. J. H., and Malhotra, G. Lost in Latent Space: Disentangled Models and the Challenge of Combinatorial Generalisation, June 2024. URL <http://arxiv.org/abs/2204.02283>. arXiv:2204.02283 [cs]. 1, 6, 8
- Morioka, H. and Hyvärinen, A. Connectivity-contrastive learning: Combining causal discovery and representation learning for multimodal data. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, pp. 3399–3426. PMLR, April 2023. URL <https://proceedings.mlr.press/v206/morioka23a.html>. ISSN: 2640-3498. 3, 4
- Morioka, H. and Hyvärinen, A. Causal Representation Learning Made Identifiable by Grouping of Observational Variables, October 2023. URL <http://arxiv.org/abs/2310.15709>. arXiv:2310.15709 [cs, stat]. 4
- Morioka, H., Hälvä, H., and Hyvärinen, A. Independent Innovation Analysis for Nonlinear Vector Autoregressive Process. *arXiv:2006.10944 [cs, stat]*, February 2021. URL <https://arxiv.org/abs/2006.10944>. arXiv: 2006.10944. 3, 4
- Morningstar, W., Bijamov, A., Duvarney, C., Friedman, L., Kalibhat, N., Liu, L., Mansfield, P., Rojas-Gomez, R., Singhal, K., Green, B., and Prakash, S. Augmentations vs Algorithms: What Works in Self-Supervised Learning, March 2024. URL <http://arxiv.org/abs/2403.05726>. arXiv:2403.05726 [cs]. 1, 2, 4, 6, 7, 8, 18
- Moschella, L., Maiorca, V., Fumero, M., Norelli, A., Locatello, F., and Rodolà, E. Relative representations enable zero-shot latent space communication, March 2023. URL <http://arxiv.org/abs/2209.15430>. arXiv:2209.15430 [cs]. 2
- Moutakanni, T., Oquab, M., Szafraniec, M., Vakalopoulou, M., and Bojanowski, P. You Don’t Need Data-Augmentation in Self-Supervised Learning, June 2024. URL <http://arxiv.org/abs/2406.09294>. arXiv:2406.09294 [cs]. 4
- Mészáros, A., Ujváry, S., Brendel, W., Reizinger, P., and Huszár, F. Rule Extrapolation in Language Modeling: A Study of Compositional Generalization on OOD Prompts. November 2024. URL <https://openreview.net/forum?id=Li2rpRZWjy>. 6
- Nakkiran, P., Kaplun, G., Kalimeris, D., Yang, T., Edelman, B. L., Zhang, F., and Barak, B. SGD on neural networks learns functions of increasing complexity. *CoRR*, abs/1905.11604, 2019. URL <http://arxiv.org/abs/1905.11604>. 18
- Norelli, A., Fumero, M., Maiorca, V., Moschella, L., Rodolà, E., and Locatello, F. ASIF: Coupled Data Turns Unimodal Models to Multimodal Without Training, February 2023. URL <http://arxiv.org/abs/2210.01738>. arXiv:2210.01738 [cs]. 2
- Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.-Y., Li, S.-W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., and Bojanowski, P. DINOv2: Learning Robust Visual Features without Supervision, February 2024. URL <http://arxiv.org/abs/2304.07193>. arXiv:2304.07193 [cs]. 6, 7, 8
- Papayan, V., Han, X., and Donoho, D. L. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020. 6

- Park, K., Choe, Y. J., and Veitch, V. The Linear Representation Hypothesis and the Geometry of Large Language Models, November 2023. URL <http://arxiv.org/abs/2311.03658>. arXiv:2311.03658 [cs, stat]. 2
- Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2 edition, 2009. ISBN 978-0-511-80316-1. doi: 10.1017/CBO9780511803161. URL <http://ebooks.cambridge.org/ref/id/CBO9780511803161>. 3
- Power, A., Burda, Y., Edwards, H., Babuschkin, I., and Misra, V. Grokking: Generalization Beyond Overfitting on Small Algorithmic Datasets, January 2022. URL <http://arxiv.org/abs/2201.02177>. arXiv:2201.02177 [cs]. 5
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language Models are Unsupervised Multi-task Learners. pp. 24, 2018. 1
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning transferable visual models from natural language supervision, 2021. URL <https://arxiv.org/abs/2103.00020>. 1, 2
- Rajendran, G., Reizinger, P., Brendel, W., and Ravikumar, P. An Interventional Perspective on Identifiability in Gaussian LTI Systems with Independent Component Analysis, November 2023. URL <http://arxiv.org/abs/2311.18048>. arXiv:2311.18048 [cs, eess, stat]. 3, 4
- Raventós, A., Paul, M., Chen, F., and Ganguli, S. Pretraining task diversity and the emergence of non-Bayesian in-context learning for regression. *Advances in Neural Information Processing Systems*, 36:14228–14246, December 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/hash/2e10b2c2e1aa4f8083c37dfe269873f8-Abstract-Conference.html](https://proceedings.neurips.cc/paper_files/paper/2023/hash/2e10b2c2e1aa4f8083c37dfe269873f8-Abstract-Conference.html). 4
- Reizinger, P., Gresele, L., Brady, J., Kügelgen, J. V., Zietlow, D., Schölkopf, B., Martius, G., Brendel, W., and Besserve, M. Embrace the Gap: VAEs Perform Independent Mechanism Analysis. October 2022. URL <https://openreview.net/forum?id=G4GpqX4bKAH>. 3
- Reizinger, P., Bizeul, A., Juhos, A., Vogt, J. E., Balestriero, R., Brendel, W., and Klindt, D. Cross-Entropy Is All You Need To Invert the Data Generating Process. October 2024a. URL <https://openreview.net/forum?id=hrqN0xpItr>. 3, 4, 6, 7, 8, 18
- Reizinger, P., Guo, S., Huszár, F., Schölkopf, B., and Brendel, W. Identifiable Exchangeable Mechanisms for Causal Structure and Representation Learning. October 2024b. URL <https://openreview.net/forum?id=k03mB41vyM>. 3
- Reizinger, P., Ujváry, S., Mészáros, A., Kerekes, A., Brendel, W., and Huszár, F. Position: Understanding LLMs Requires More Than Statistical Generalization. June 2024c. URL <https://openreview.net/forum?id=pVyOchWUBa>. 5, 18
- Roeder, G., Metz, L., and Kingma, D. P. On Linear Identifiability of Learned Representations. *arXiv:2007.00810 [cs, stat]*, July 2020. URL <http://arxiv.org/abs/2007.00810>. arXiv: 2007.00810. 2, 3, 5, 6, 18
- Ruoss, A., Delétang, G., Genewein, T., Grau-Moya, J., Csordás, R., Bennani, M., Legg, S., and Veness, J. Randomized Positional Encodings Boost Length Generalization of Transformers, May 2023. URL <http://arxiv.org/abs/2305.16843>. arXiv:2305.16843 [cs, stat]. 6
- Rusak, E., Reizinger, P., Juhos, A., Bringmann, O., Zimmermann, R. S., and Brendel, W. InfoNCE: Identifying the Gap Between Theory and Practice, June 2024. URL <http://arxiv.org/abs/2407.00143>. arXiv:2407.00143 [cs, stat]. 3, 4, 5, 7, 18
- Saunshi, N., Malladi, S., and Arora, S. A Mathematical Exploration of Why Language Models Help Solve Downstream Tasks, April 2021. URL <http://arxiv.org/abs/2010.03648>. arXiv:2010.03648 [cs, stat]. 6
- Schneider, S., Lee, J. H., and Mathis, M. W. Learnable latent embeddings for joint behavioural and neural analysis. *Nature*, 617(7960):360–368, 2023. 4
- Schott, L., von Kügelgen, J., Träuble, F., Gehler, P., Russell, C., Bethge, M., Schölkopf, B., Locatello, F., and Brendel, W. Visual Representation Learning Does Not Generalize Strongly Within the Same Domain. *arXiv:2107.08221 [cs]*, July 2021. URL <http://arxiv.org/abs/2107.08221>. arXiv: 2107.08221. 1, 6, 8
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., and Bengio, Y. Towards Causal Representation Learning. *arXiv:2102.11107 [cs]*, February 2021. URL <http://arxiv.org/abs/2102.11107>. arXiv: 2102.11107 version: 1. 3
- Shimizu, S., Hoyer, P. O., Hyvarinen, A., and Kerminen, A. A Linear Non-Gaussian Acyclic Model for Causal Discovery. pp. 28, 2006. 18

- Shwartz-Ziv, R., Balestrieri, R., and LeCun, Y. What Do We Maximize in Self-Supervised Learning?, July 2022. URL <http://arxiv.org/abs/2207.10081>. arXiv:2207.10081 [cs]. 7
- Simon, J. B., Knutins, M., Ziyin, L., Geisz, D., Fetterman, A. J., and Albrecht, J. On the Stepwise Nature of Self-Supervised Learning, May 2023. URL <http://arxiv.org/abs/2303.15438>. arXiv:2303.15438 [cs]. 5, 7, 18
- Song, Z., Su, X., Wang, J., Qiang, W., Zheng, C., and Sun, F. Towards the Sparseness of Projection Head in Self-Supervised Learning, July 2023. URL <http://arxiv.org/abs/2307.08913>. arXiv:2307.08913 [cs]. 6
- Sorscher, B., Geirhos, R., Shekhar, S., Ganguli, S., and Morcos, A. Beyond neural scaling laws: beating power law scaling via data pruning. *Advances in Neural Information Processing Systems*, 35:19523–19536, December 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/hash/7b75da9b61eda40fa35453ee5d077df6-Abstract-Conference.html](https://proceedings.neurips.cc/paper_files/paper/2022/hash/7b75da9b61eda40fa35453ee5d077df6-Abstract-Conference.html). 8
- 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. 3
- Sutton, R. The bitter lesson. *Incomplete Ideas (blog)*, 13(1): 38, 2019. 3, 8
- Tenenbaum, J. B., Silva, V. d., and Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, December 2000. ISSN 1095-9203. doi: 10.1126/science.290.5500.2319. URL <http://dx.doi.org/10.1126/science.290.5500.2319>. 7
- Thilak, V., Huang, C., Saremi, O., Dinh, L., Goh, H., Nakkiran, P., Susskind, J. M., and Littwin, E. Lidar: Sensing linear probing performance in joint embedding ssl architectures, 2023. URL <https://arxiv.org/abs/2312.04000>. 8
- Tian, Y. Understanding the Role of Nonlinearity in Training Dynamics of Contrastive Learning. November 2022. URL <https://openreview.net/forum?id=s130rTE3UX>. 6
- von Kügelgen, J. Identifiable Causal Representation Learning: Unsupervised, Multi-View, and Multi-Environment. March 2024. URL <https://www.repository.cam.ac.uk/handle/1810/365627>. 3
- von Kügelgen, J., Sharma, Y., Gresele, L., Brendel, W., Schölkopf, B., Besserve, M., and Locatello, F. Self-Supervised Learning with Data Augmentations Provably Isolates Content from Style, June 2021. URL <http://arxiv.org/abs/2106.04619>. arXiv: 2106.04619. 2, 4, 5, 18
- von Kügelgen, J., Besserve, M., Wendong, L., Gresele, L., Kekić, A., Bareinboim, E., Blei, D. M., and Schölkopf, B. Nonparametric Identifiability of Causal Representations from Unknown Interventions, October 2023. URL <http://arxiv.org/abs/2306.00542>. arXiv:2306.00542 [cs, stat]. 3
- Wang, T. and Isola, P. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. arXiv:2005.10242 [cs, stat], November 2020. URL <http://arxiv.org/abs/2005.10242>. arXiv: 2005.10242. 1, 3, 7
- Wang, Y., Albrecht, C. M., Braham, N. A. A., Mou, L., and Zhu, X. X. Self-supervised learning in remote sensing: A review. *IEEE Geoscience and Remote Sensing Magazine*, 10(4):213–247, 2022. 2
- Watanabe, S. *Algebraic Geometry and Statistical Learning Theory*. Cambridge University Press, 1 edition, August 2009. ISBN 978-0-521-86467-1 978-0-511-80047-4. doi: 10.1017/CBO9780511800474. URL <https://www.cambridge.org/core/product/identifier/9780511800474/type/book>. 2, 4
- Watanabe, S. *Mathematical theory of Bayesian statistics*. CRC Press, Taylor & Francis Group, Boca Raton London New York, first issued in paperback edition, 2020. ISBN 978-1-4822-3806-8 978-0-367-73481-7. 2, 4
- Wendong, L., Kekić, A., von Kügelgen, J., Buchholz, S., Besserve, M., Gresele, L., and Schölkopf, B. Causal Component Analysis, October 2023. URL <http://arxiv.org/abs/2305.17225>. arXiv:2305.17225 [cs, stat]. 3
- Wiedemer, T., Brady, J., Panfilov, A., Juhos, A., Bethge, M., and Brendel, W. Provable Compositional Generalization for Object-Centric Learning, October 2023a. URL <http://arxiv.org/abs/2310.05327>. arXiv:2310.05327 [cs]. 4, 6, 8, 18
- Wiedemer, T., Mayilvahanan, P., Bethge, M., and Brendel, W. Compositional Generalization from First Principles, July 2023b. URL <http://arxiv.org/abs/2307.05596>. arXiv:2307.05596 [cs, stat]. 4, 6, 8, 18



- Xie, S. M., Raghunathan, A., Liang, P., and Ma, T. An Explanation of In-context Learning as Implicit Bayesian Inference, July 2022. URL <http://arxiv.org/abs/2111.02080>. arXiv:2111.02080 [cs]. 6
- Xue, Y., Gan, E., Ni, J., Joshi, S., and Mirzasoleiman, B. Investigating the Benefits of Projection Head for Representation Learning, March 2024. URL <http://arxiv.org/abs/2403.11391>. arXiv:2403.11391 [cs]. 6
- Yao, D., Muller, C., and Locatello, F. Marrying Causal Representation Learning with Dynamical Systems for Science, May 2024. URL <http://arxiv.org/abs/2405.13888>. arXiv:2405.13888 [cs, stat]. 4
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. *arXiv:2103.03230 [cs, q-bio]*, June 2021. URL <http://arxiv.org/abs/2103.03230>. arXiv: 2103.03230. 6
- Zhai, X., Kolesnikov, A., Houlsby, N., and Beyer, L. Scaling Vision Transformers, June 2022. URL <http://arxiv.org/abs/2106.04560>. arXiv:2106.04560 [cs]. 8
- Zhai, X., Mustafa, B., Kolesnikov, A., and Beyer, L. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 11975–11986, 2023. 2
- Zhou, D. and Wei, X.-X. Learning identifiable and interpretable latent models of high-dimensional neural activity using pi-vae. *Advances in Neural Information Processing Systems*, 33:7234–7247, 2020. 4
- Zimmermann, R. S., Sharma, Y., Schneider, S., Bethge, M., and Brendel, W. Contrastive Learning Inverts the Data Generating Process. *arXiv:2102.08850 [cs]*, February 2021. URL <http://arxiv.org/abs/2102.08850>. arXiv: 2102.08850. 3, 4, 7, 8, 18

## A. On absolute vs relative identifiability

We need to tackle one nuance when talking about identifiability and the alignment of representations: namely, what do we mean by identifiability? There are two notions in the literature, which we will call *relative* and *absolute identifiability*. These are defined as follows:

**Definition A.1** (Absolute identifiability of the ground-truth representation). Given a ground-truth probabilistic model  $p_\theta$ , and a learned family of probabilistic models parametrized by a neural network at the optimum of its training loss  $p_{\hat{\theta}} : \hat{\theta} \in \hat{\Theta}$ , the inferred model identifies the ground-truth representation if  $\forall \hat{\theta} \in \hat{\Theta}$  and a particular equivalence relationship  $\sim_{abs}$  it holds that  $p_{\hat{\theta}} \sim_{abs} p_\theta$ .

That is, absolute identifiability means that there is an equivalence relationship (e.g., permutation or scaling) between the learned (family of) representations and the ground-truth one. “Family” refers to, e.g., training the same model with different seeds. This identifiability notion is more prevalent in the ICA literature (Hyvärinen & Pajunen, 1999; Hyvärinen, 2013; Hyvärinen et al., 2023).

On the other hand, relative identifiability posits an equivalence relationship between two trained models, i.e., it does not make any claims about the ground-truth representation.

**Definition A.2** (Relative identifiability of a pair of learned representations). Given a pair of learned probabilistic model families  $p_{\theta_1}, p_{\theta_2} : \theta_1 \in \Theta_1, \theta_2 \in \Theta_2$  parametrized by (not necessarily the same) neural networks at the optimum of their respective training losses, they are said to be identified in a relative sense if there exists an equivalence relationship  $\sim_{rel}$  such that  $\forall \theta_1 \in \Theta_1, \theta_2 \in \Theta_2 : p_{\theta_1} \sim_{rel} p_{\theta_2}$ .

## B. The Platonic Representation Hypothesis revisited

Huh et al. (2024) argued that the i) function class, ii) dataset, and iii) regularization, determine the convergence of representations, which are important in SSL (Cabannes et al., 2023; Balestriero et al., 2023; Morningstar et al., 2024). By synthesizing evidence from the identifiability literature, we showcase the mathematical results answering some of these questions. However, we will also show the negative results and the limits of current IT, enabling us to precisely specify future research directions. Before that, we emphasize that the PRH’s claims are about the similarity of two *learned* representations (such as in Roeder et al. (2020)), whereas most identifiability results are about an *absolute* sense between a learned model and the assumed underlying DGP—we discuss the differences in Appx. A.

**Function class and model capacity.** Restricting the function class is a conventional technique to prove identifiability (Hyvärinen & Pajunen, 1999; Locatello et al., 2019; Bloem-Reddy & Teh, 2020; Buchholz et al., 2022). That is, in some cases no restriction can imply non-identifiability, indicated by the impossibility results of Locatello et al. (2019); Hyvärinen & Pajunen (1999). A restriction of the functions class in form of a log-linear model was also shown to be key for recent identifiability results in (self-)supervised learning (Hyvärinen & Morioka, 2016; Hyvärinen et al., 2019; Roeder et al., 2020; Reizinger et al., 2024a; Shimizu et al., 2006; Hoyer et al., 2008; Lachapelle et al., 2020; Gresele et al., 2021). Compositional generalization also hinges on particular (mainly additive) decoder structures (Lachapelle et al., 2023; Brady et al., 2023; Wiedemer et al., 2023a;b; Brady et al., 2025; Mahajan et al., 2024; Ahuja et al., 2022). Such results indicate that the PRH’s validity is potentially restricted to a specific model class. The investigation of Ciernik et al. (2024) also suggests a varying degree of similarity even between self-supervised models (Fig. 1).

**Data set and training objective.** Many self-supervised models optimize an objective that corresponds to the cross entropy between the distributions of the underlying and the learned DGPs (Zimmermann et al., 2021; Huh et al., 2024; Reizinger et al., 2024a). This implies that, e.g., changing the way positive pairs are generated implies adapting the loss function to avoid a mismatch (Zimmermann et al., 2021; Rusak et al., 2024). von Kügelgen et al. (2021) proved that when some latent factors are not changed by the augmentations, then only those are provably identifiable—others collapse, as observed in the projector phenomenon (Chen et al., 2020; Rusak et al., 2024; Jing et al., 2022). This supports empirical observations about the importance of augmentations (Morningstar et al., 2024), and suggests that different sets of augmentations can have different “Platonic ideals”.

**Regularization.** Huh et al. (2024) posited that deep networks are biased towards learning simple functions (Nakkiran et al., 2019). In a discussion about the non-identifiability of LLMs, Reizinger et al. (2024c) noted that the theoretical understanding of a wide range of inductive biases is still missing—this applies to SSL in general, too—at least for inductive biases beyond the ones expressed via the loss function, the function class, or data augmentations. Rusak et al. (2024) suggested an anisotropic conditional can worsen the learning dynamics of high-variance factors, though more work is required to precisely understand nonlinear learning dynamics in SSL (Simon et al., 2023).

**Summary.** In our view, the PRH made a first important step formalizing the similarities of learned representations. However, it does not provide a clear theoretical framework to precisely investigate such questions

## **C. Acronyms**

**BSS** Blind Source Separation

**CL** Contrastive Learning

**CRL** Causal Representation Learning

**DGP** data generating process

**i.i.d.** independent and identically distributed

**ICA** Independent Component Analysis

**IT** identifiability theory

**LLM** Large Language Model

**LT** learning theory

**OOD** out-of-distribution

**PRH** Platonic Representation Hypothesis

**SITh** Singular Identifiability Theory

**SLT** Singular Learning Theory

**SSL** self-supervised learning

**vMF** von Mises-Fisher