

---

# Synergies between Disentanglement and Sparsity: Generalization and Identifiability in Multi-Task Learning

---

Sébastien Lachapelle<sup>\*1</sup> Tristan Deleu<sup>\*1</sup> Divyat Mahajan<sup>1</sup> Ioannis Mitliagkas<sup>1,2</sup> Yoshua Bengio<sup>1,2</sup>  
Simon Lacoste-Julien<sup>1,2</sup> Quentin Bertrand<sup>1</sup>

## Abstract

Although disentangled representations are often said to be beneficial for downstream tasks, current empirical and theoretical understanding is limited. In this work, we provide evidence that disentangled representations coupled with sparse task-specific predictors improve generalization. In the context of multi-task learning, we prove a new identifiability result that provides conditions under which maximally sparse predictors yield disentangled representations. Motivated by this theoretical result, we propose a practical approach to learn disentangled representations based on a sparsity-promoting bi-level optimization problem. Finally, we explore a meta-learning version of this algorithm based on group Lasso multiclass SVM predictors, for which we derive a tractable dual formulation. It obtains competitive results on standard few-shot classification benchmarks, while each task is using only a fraction of the learned representations.

## 1. Introduction

The recent literature on self-supervised learning has provided evidence that learning a representation on large corpora of data can yield strong performances on a wide variety of downstream tasks (Devlin et al., 2018; Chen et al., 2020), especially in few-shot learning scenarios where the training data for these tasks is limited (Brown et al., 2020b; Dosovitskiy et al., 2021; Radford et al., 2021). Beyond transferring across multiple tasks, these learned representations also lead to improved robustness against distribution shifts (Wortsman et al., 2022) as well as stunning text-

conditioned image generation (Ramesh et al., 2022). However, preliminary assessments of the latter have highlighted shortcomings related to compositionality (Marcus et al., 2022), suggesting new algorithmic innovations are needed.

Another line of work has argued for the integration of ideas from causality to make progress towards more robust and transferable machine learning systems (Pearl, 2019; Schölkopf, 2019; Goyal & Bengio, 2022). *Causal representation learning* has emerged recently as a field aiming to define and learn representations suited for causal reasoning (Schölkopf et al., 2021). This set of ideas is strongly related to learning *disentangled representations* (Bengio et al., 2013). Informally, a representation is considered disentangled when its components are in one-to-one correspondence with natural and interpretable factors of variations, such as object positions, colors or shapes. Although a plethora of works have investigated theoretically under which conditions disentanglement is possible through the lens of identifiability (Hyvärinen & Morioka, 2016; 2017; Hyvärinen et al., 2019; Khemakhem et al., 2020a; Locatello et al., 2020a; Klindt et al., 2021; Von Kügelgen et al., 2021; Gresele et al., 2021; Lachapelle et al., 2022; Lippe et al., 2022b; Ahuja et al., 2022c), fewer works have tackled *how a disentangled representation could be beneficial for downstream tasks*. Those who did mainly provide empirical rather than theoretical evidence for or against its usefulness (Locatello et al., 2019; van Steenkiste et al., 2019; Miladinovi et al., 2019; Dittadi et al., 2021; Montero et al., 2021). We believe our work can bring some theoretical insights as to when and why disentanglement can help.

In this work, we explore synergies between disentanglement and sparse task-specific predictors in the context of multi-task learning. At the heart of our contributions is the assumption that only a small subset of all factors of variations are useful for each downstream task, and this subset might change from one task to another. We will refer to such tasks as *sparse tasks*, and their corresponding sets of useful factors as their *supports*. This assumption was initially suggested by Bengio et al. (2013, Section 3.5): “the feature set being trained may be destined to be used in multiple tasks that may have distinct [and unknown] sub-

---

<sup>\*</sup>Equal contribution <sup>1</sup>Mila & DIRO, Université de Montréal <sup>2</sup>Canada CIFAR AI Chair. Correspondence to: Sébastien Lachapelle <lachaseb@mila.quebec>, Tristan Deleu <deleutri@mila.quebec>.

sets of relevant features. Considerations such as these lead us to the conclusion that the most robust approach to feature learning is to disentangle as many factors as possible, discarding as little information about the data as is practical”. This strategy is in line with the current self-supervised learning trend (Radford et al., 2021), except for its focus on disentanglement.

### 1.1. Contributions

1. We formalize this “sparse task assumption” and argue theoretically and empirically how, when it holds, a disentangled representation coupled with a sparsity-regularized task-specific predictor can generalize better than their entangled counterparts (Section 2).
2. We introduce a novel identifiability result (Theorem 3.1) which shows how one can leverage multiple sparse supervised tasks to learn a shared disentangled representation by regularizing the task-specific predictors to be maximally sparse (Section 3.2). We note that the usage of supervision is in line with many recent results which leverages more or less weak forms of supervision to guarantee identifiability. Contrary to many existing identifiability results, ours allows for statistically dependent latent factors and a non-invertible map between observations and latents.
3. Motivated by this result, we propose a tractable bi-level optimization (Problem (6)) to learn the shared representation while regularizing the task-specific predictors to be sparse (Section 3.4). We validate our theory by showing our approach can indeed disentangle latent factors on tasks constructed from the 3D Shapes dataset (Burgess & Kim, 2018).
4. Finally, we draw a connection between this bi-level optimization problem and formulations from the meta-learning literature. Inspired by our identifiability result, we enhance an existing method (Lee et al., 2019), where the task-specific predictors are now group-sparse SVMs. We show that this new meta-learning algorithm achieves competitive performance on the *mini*ImageNet benchmark (Vinyals et al., 2016), while only using a fraction of the representation.

We emphasize that, although related, the theoretical contributions of Sections 2 & 3 are distinct and stand of their own. Indeed, Section 2 shows how disentangled representations combined with sparsity regularization can improve generalization, while Section 3 shows how regularizing task-specific predictors to be sparse can induce disentanglement in a multi-task learning setting.

### 1.2. Background

We start by introducing formally the notion of entangled and disentangled representations.

First, we assume the existence of some ground-truth encoder function  $f_{\theta} : \mathbb{R}^d \rightarrow \mathbb{R}^m$  that maps observations  $x \in \mathcal{X} \subseteq \mathbb{R}^d$ , e.g., images, to its corresponding interpretable and usually lower dimensional representation  $f_{\theta}(x) \in \mathbb{R}^m$ ,  $m \leq d$ . The exact form of this ground-truth encoder depends on the task at hand, but also on what the machine learning practitioner considers as interpretable. The learned encoder function is denoted by  $f_{\hat{\theta}} : \mathbb{R}^d \rightarrow \mathbb{R}^m$ , and should not be conflated with the ground-truth representation  $f_{\theta}$ . For example,  $f_{\hat{\theta}}$  can be parametrized by a neural network. Throughout, we are going to use the following definition of disentanglement.

**Definition 1.1** (Disentangled Representation, Khemakhem et al. 2020a; Lachapelle et al. 2022). A learned encoder function  $f_{\hat{\theta}} : \mathbb{R}^d \rightarrow \mathbb{R}^m$  is said to be *disentangled w.r.t. the ground-truth representation  $f_{\theta}$*  when there exists an invertible diagonal matrix  $D$  and a permutation matrix  $P$  such that, for all  $x \in \mathcal{X}$ ,  $f_{\hat{\theta}}(x) = DPf_{\theta}(x)$ .

Intuitively, a representation is disentangled when there is a one-to-one correspondence between its components and those of the ground-truth representation, up to rescaling. When an encoder  $f_{\hat{\theta}}$  is not disentangled, we say it is *entangled*. Note that there exist less stringent notions of disentanglement which allow for component-wise nonlinear invertible transformations of the factors (Hyvärinen & Morioka, 2017; Hyvärinen et al., 2019).

**Notation.** Capital bold letters denote matrices and lowercase bold letters denote vectors. The set of integers from 1 to  $n$  is denoted by  $[n]$ . We write  $\|\cdot\|$  for the Euclidean norm on vectors and the Frobenius norm on matrices. For a matrix  $A \in \mathbb{R}^{k \times m}$ ,  $\|A\|_{2,1} = \sum_{j=1}^m \|A_{:j}\|$ , and  $\|A\|_{2,0} = \sum_{j=1}^m \mathbb{1}_{\|A_{:j}\| \neq 0}$ , where  $\mathbb{1}$  is the indicator function. The ground-truth parameter of the encoder function is  $\theta$ , while that of the learned representation is  $\hat{\theta}$ . We follow this convention for all the parameters throughout. Table 1 in Appendix A summarizes all the notation.

## 2. Disentanglement and Sparse Task-Specific Predictors Improve Generalization

In this section, we show that for any *linearly equivalent* representation (entangled or disentangled), the maximum likelihood estimator defined in Problem (1) yields the same model (Proposition 2.2). However, we also show that disentangled representations have better generalization properties when the task-specific predictor is regularized to be sparse. (Proposition 2.4 and Fig. 1). Our analysis is centred around the following assumption.

**Assumption 2.1** (Linear equivalence). The learned encoder  $f_{\hat{\theta}}$  is *linearly equivalent* to the ground-truth encoder  $f_{\theta}$ , i.e., there exists an invertible matrix  $L$  such that, for all  $x \in \mathcal{X}$ ,  $f_{\hat{\theta}}(x) = Lf_{\theta}(x)$ .

Note that similar notions of linear equivalence were used e.g. by Hyvärinen et al. (2019); Khemakhem et al. (2020a); Roeder et al. (2021)

Despite being assumed linearly equivalent, the learned representation  $\mathbf{f}_{\hat{\theta}}$  might not be disentangled (Definition 1.1); in that case, we say the representation is *linearly entangled*. When we refer to a disentangled representation, we write  $\mathbf{L} := \mathbf{D}\mathbf{P}$ . Roeder et al. (2021) have shown that many common methods learn representations identifiable up to linear equivalence, such as deep neural networks for classification, contrastive learning (Oord et al., 2018; Radford et al., 2021) and autoregressive language models (Mikolov et al., 2010; Brown et al., 2020a).

## 2.1. MLE invariance to linear feature transformations

Consider the following maximum likelihood estimator (MLE):<sup>1</sup>

$$\hat{\mathbf{W}}_n^{(\hat{\theta})} := \arg \max_{\tilde{\mathbf{W}}} \sum_{(\mathbf{x}, y) \in \mathcal{D}} \log p(y; \boldsymbol{\eta} = \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})), \quad (1)$$

where  $y$  denotes the label,  $\mathcal{D} := \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=1}^n$  is the dataset,  $p(y; \boldsymbol{\eta})$  is a distribution over labels<sup>2</sup> parameterized by  $\boldsymbol{\eta} \in \mathbb{R}^k$ , and  $\tilde{\mathbf{W}} \in \mathbb{R}^{k \times m}$  is the *task-specific predictor*. The following result shows that the model estimated via maximum likelihood defined in Problem (1) is invariant to invertible linear transformations of the features. Note that it is an almost direct consequence of the invariance of MLE to reparametrization (Casella & Berger, 2001, Thm. 7.2.10). See Appendix A for a proof.

**Proposition 2.2.** *Let  $\hat{\mathbf{W}}_n^{(\hat{\theta})}$  and  $\hat{\mathbf{W}}_n^{(\theta)}$  be the solutions to Problem (1) with the representations  $\mathbf{f}_{\hat{\theta}}$  and  $\mathbf{f}_{\theta}$ , respectively (which we assume are unique). If  $\mathbf{f}_{\hat{\theta}}$  and  $\mathbf{f}_{\theta}$  are linearly equivalent (Assumption 2.1), then we have,  $\forall \mathbf{x} \in \mathcal{X}$ ,  $\hat{\mathbf{W}}_n^{(\hat{\theta})} \mathbf{f}_{\hat{\theta}}(\mathbf{x}) = \hat{\mathbf{W}}_n^{(\theta)} \mathbf{f}_{\theta}(\mathbf{x})$ .*

Proposition 2.2 shows that the model  $p(y; \hat{\mathbf{W}}_n^{(\hat{\theta})} \mathbf{f}_{\hat{\theta}}(\mathbf{x}))$  learned by Problem (1) is independent of  $\mathbf{L}$ , i.e., *the learned model is the same for disentangled and linearly entangled representations*. We thus expect both disentangled and linearly entangled representations to perform identically on downstream tasks.

## 2.2. An advantage of disentangled representations

We are now going to see how adding sparsity regularization to Problem (1) favors the disentangled representation when the ground-truth data generating process is truly sparse.

**Assumption 2.3** (Data generation process). The input-label pairs are i.i.d. samples from the distribution

$p(\mathbf{x}, y) := p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x}))p(\mathbf{x})$ , where  $\mathbf{W} \in \mathbb{R}^{k \times m}$  is the ground-truth coefficient matrix such that  $\|\mathbf{W}\|_{2,0} = \ell$ .

To formalize the hypothesis that *only a subset of the features  $\mathbf{f}_{\theta}(\mathbf{x})$  are actually useful to predict the target  $y$* , we assume that the ground-truth coefficient matrix  $\mathbf{W}$  is column sparse, i.e.,  $\|\tilde{\mathbf{W}}\|_{2,0} = \ell < m$ . Under this assumption, it is natural to constrain the MLE as such:

$$\hat{\mathbf{W}}_n^{(\hat{\theta}, \ell)} := \arg \max_{\|\tilde{\mathbf{W}}\|_{2,0} \leq \ell} \sum_{(\mathbf{x}, y) \in \mathcal{D}} \log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \quad (2)$$

To analyze the impact of this additional constraint on the generalization error, we consider both the estimation error (a.k.a. variance) and the approximation error (a.k.a. bias) separately (Mohri et al., 2018, Chapter 4).

**Estimation error.** The sparsity constraint of Problem (2) decreases the size of the hypothesis class considered to minimize the negative log-likelihood and should thus yield a decrease in estimation error for both entangled and disentangled representations (i.e., reduce overfitting). Sparsity regularization is a well-understood approach to control the complexity of a predictor, see for example Bickel et al. (2009); Lounici et al. (2011a); Mohri et al. (2018).

**Approximation error.** Disentangled and entangled representations differ in how the sparsity constraint of Problem (2) impacts their approximation errors. The following proposition will help us see how this regularization favors disentangled representations over entangled ones.

**Proposition 2.4.** *Let  $\hat{\mathbf{W}}_{\infty}^{(\hat{\theta})}$  be the (assumed unique) solution of the population-based MLE,  $\arg \max_{\tilde{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x}))$ . If Assumption 2.1 (linear equivalence) & Assumption 2.3 (data generating process) hold,  $\hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} = \mathbf{W} \mathbf{L}^{-1}$ .*

From Proposition 2.4, one can see that if the representation  $\mathbf{f}_{\hat{\theta}}$  is disentangled ( $\mathbf{L} = \mathbf{D}\mathbf{P}$ ), then

$$\|\hat{\mathbf{W}}_{\infty}^{(\hat{\theta})}\|_{2,0} = \|\mathbf{W}(\mathbf{D}\mathbf{P})^{-1}\|_{2,0} = \|\mathbf{W}\|_{2,0} = \ell.$$

Thus, the sparsity constraint in Problem (2) does not exclude the population MLE estimator from its hypothesis class which means no approximation error is entailed (no bias). Contrarily, when  $\mathbf{f}_{\hat{\theta}}$  is linearly entangled, the population MLE might have more nonzero columns than the ground-truth (since  $\mathbf{L}^{-1}$  might destroy the sparsity of  $\mathbf{W}$ ), and thus would be excluded from the hypothesis space of Problem (2), which means an approximation error is introduced.

**Conclusion.** The above points suggest that *if the ground-truth task is sufficiently sparse, the disentangled representation should benefit from sparsity regularization (assuming*

<sup>1</sup>We assume the solution is unique.

<sup>2</sup> $p(y; \boldsymbol{\eta})$  could be a Gaussian density (regression) or a categorical distribution (classification).

the number of samples is low) because it reduces the estimation error (variance) without increasing the approximation error (bias). In contrast, an entangled representation might not benefit from sparsity regularization if the increase in approximation error is more important than the reduction in estimation error.

**Empirical validation (Fig. 1).** We now present a simple simulated experiment that illustrates the above claim that *disentangled representations coupled with sparsity regularization can yield better generalization*. Fig. 1 compares the generalization performances of  $L_1$  and  $L_2$ -penalized linear regressions (Tibshirani, 1996; Hoerl & Kennard, 1970), computed on the top of both disentangled and linearly entangled representations, which are frozen during training.  $L_1$ -penalized linear regression coupled with the disentangled representation yields better generalization than other alternatives when  $\ell/m = 5\%$  and when the number of samples is very small. One can also see that disentanglement, sparsity regularization, and sufficient sparsity in the ground-truth data generating process are necessary for significant improvements, in line with our discussion. Lastly, all methods yield similar performance when the number of samples grows. More details and discussions can be found in Appendix D.1.

### 3. Sparse Multi-Task Learning for Disentanglement

In Section 2, we argued that disentangled representations can improve generalization when combined with sparse task-specific predictors, but we did not mention how to obtain a disentangled representation in the first place. In this section, we first provide a new identification result (Theorem 3.1, Section 3.2), which states that in the multi-task learning setting, regularizing the task-specific predictors to be sparse can yield disentangled representations. Then, in Section 3.4, we provide a practical way to learn disentangled representations motivated by our identifiability result.

#### 3.1. Task & data generating process

Throughout this section, we assume the learner is given a set of  $T$  datasets  $\{\mathcal{D}_1, \dots, \mathcal{D}_T\}$  where each dataset  $\mathcal{D}_t := \{(\mathbf{x}^{(t,i)}, y^{(t,i)})\}_{i=1}^n$  consists of  $n$  couples of input  $\mathbf{x} \in \mathbb{R}^d$  and label  $y \in \mathcal{Y}$ . The set of labels  $\mathcal{Y}$  might contain either class indices or real values, depending on whether we are concerned with classification or regression tasks.

Our theory relies on the assumption that, for each task  $t$ , the dataset  $\mathcal{D}_t$  is made of i.i.d. samples from the distribution

$$p(\mathbf{x}, y | \mathbf{W}^{(t)}) := p(y; \mathbf{W}^{(t)} \mathbf{f}_\theta(\mathbf{x}))p(\mathbf{x} | \mathbf{W}^{(t)}), \quad (3)$$

where  $\mathbf{W}^{(t)} \in \mathbb{R}^{k \times m}$  is the task-specific ground-truth coefficient matrix. We emphasize that the representation  $\mathbf{f}_\theta$




Figure 1. Test performance for the entangled and disentangled representation using Lasso and Ridge regression. All the results are averaged over 10 seeds, with standard error shown in error bars.

is shared across all the tasks while the coefficient matrices  $\mathbf{W}^{(t)}$  are task-specific. Also note that the distribution over  $\mathbf{x}$  is allowed to change from one task to another. However, we assume that its support,  $\mathcal{X}$ , is fixed across tasks.

We further assume that the task-specific matrices  $\mathbf{W}^{(t)}$  are i.i.d. samples from some probability measure  $\mathbb{P}_{\mathbf{W}}$  with support  $\mathcal{W}$ . We will see in Section 3.3 that the most critical assumptions of our theory concern  $\mathbb{P}_{\mathbf{W}}$ .

#### 3.2. Main identifiability result

We are now ready to show the main theoretical result of this work, which provides a bi-level optimization problem for which the optimal representations are guaranteed to be disentangled. It assumes infinitely many tasks are observed, with task-specific ground-truth matrices  $\mathbf{W}$  sampled from  $\mathbb{P}_{\mathbf{W}}$ . We denote by  $\hat{\mathbf{W}}^{(\mathbf{W})}$  the task-specific estimator of  $\mathbf{W}$ . We delay the presentation of its technical assumptions to Section 3.3. See Appendix B.2 for a proof.

**Theorem 3.1** (Sparse multi-task learning for disentanglement). *Let  $\hat{\theta}$  be a minimizer of*

$$\begin{aligned} \min_{\hat{\theta}} \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \quad (4) \\ \text{s.t. } \hat{\mathbf{W}}^{(\mathbf{W})} \in \arg \min_{\tilde{\mathbf{W}} \text{ s.t. } \|\tilde{\mathbf{W}}\|_{2,0} \leq \|\mathbf{W}\|_{2,0}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})), \end{aligned}$$

where the constraint holds for all  $\mathbf{W} \in \mathcal{W}$  and where  $\mathbb{P}_{\mathbf{W}}$  and  $p(\mathbf{x}, y | \mathbf{W})$  are described in Section 3.1. Under Assumptions 3.2 to 3.6 and if  $\mathbf{f}_{\hat{\theta}}$  is continuous for all  $\hat{\theta}$ ,  $\mathbf{f}_{\hat{\theta}}$  is disentangled w.r.t.  $\mathbf{f}_\theta$  (Definition 1.1).

Intuitively, this optimization problem effectively selects a representation  $\mathbf{f}_{\hat{\theta}}$  that (i) allows a perfect fit of the data distribution, and (ii) allows the task-specific estimators  $\hat{\mathbf{W}}^{(\mathbf{W})}$  to be as sparse as the ground-truth  $\mathbf{W}$ . The theorem guarantees that such a representation must be disentangled.

Under the same assumptions and with the same disentanglement guarantees, Theorem B.6 in Appendix B presents a variation of Problem (4) which enforces the weaker constraint  $\mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\mathbf{W}\|_{2,0}$ , instead of  $\|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \|\mathbf{W}\|_{2,0}$  for each task  $\mathbf{W}$  individually.




Figure 2. Illustration of Assumption 3.5 showing three examples of distribution  $\mathbb{P}_{\mathbf{W}|S}$ . The red distribution satisfies the assumption, but the blue and orange distributions do not. The red lines are level sets of a Gaussian distribution with full rank covariance. The blue line represents the support of a Gaussian distribution with a low-rank covariance. The orange dots represent a distribution with finite support. The green vector  $\mathbf{a}$  shows that the condition is violated for both the blue and the orange distribution, since, in both cases,  $\mathbf{W}_{1,S}\mathbf{a}$  and  $\mathbf{a}$  are orthogonal ( $\mathbf{W}_{1,S}\mathbf{a} = 0$ ) with probability greater than zero.

**Characteristic features of our theory.** (i) Contrary to most identifiability results for disentanglement (Section 4), we do not assume the observations  $\mathbf{x}$  are generated by transforming a latent random vector  $\mathbf{z}$  through a bijective decoder  $\mathbf{g}$ . Instead, we assume the existence of a **not necessarily invertible ground-truth feature extractor**  $\mathbf{f}_\theta(\mathbf{x})$  from which the labels can be predicted using only a subset of its components in every task. (ii) Most previous works make assumptions about the distribution of latent factors, e.g., (conditional) independence, exponential family or other parametric assumptions. In contrast, we make no such assumption except a rather weak assumption on the support of the ground-truth features (Assumption 3.3). Crucially, this allows for **statistically dependent latent factors**, which we explore empirically in Section 5.1.

### 3.3. Assumptions of Theorem 3.1

We now present the technical assumptions of Theorem 3.1.

Perhaps unsurprisingly, the parameters  $\boldsymbol{\eta}$  have to be identifiable from  $p(y; \boldsymbol{\eta})$  in order for  $\mathbf{f}_\theta$  to be identifiable.

**Assumption 3.2** (Identifiability of  $\boldsymbol{\eta}$  from  $p(y; \boldsymbol{\eta})$ ).  $\text{KL}(p(y; \boldsymbol{\eta}) \parallel p(y; \tilde{\boldsymbol{\eta}})) = 0 \implies \boldsymbol{\eta} = \tilde{\boldsymbol{\eta}}$ , where KL denotes the Kullback-Leibler divergence.

This property holds, e.g., when  $p(y; \boldsymbol{\eta})$  is a Gaussian in the usual  $\mu, \sigma^2$  parameterization. Generally, it also holds for minimal parameterizations of exponential families (Wainwright & Jordan, 2008).

The following assumption requires the ground-truth representation  $\mathbf{f}_\theta(\mathbf{x})$  to vary enough such that its image cannot be trapped inside a proper subspace.

**Assumption 3.3** (Sufficient representation variability).




Figure 3. The leftmost figure represents  $\mathcal{S}$ , the set of task supports observed under the ground-truth distribution  $p(S)$ . The other figures form a verification that Assumption 3.6 holds for  $\mathcal{S}$ .

There exists  $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)} \in \mathcal{X}$  such that the matrix  $\mathbf{F} := [\mathbf{f}_\theta(\mathbf{x}^{(1)}), \dots, \mathbf{f}_\theta(\mathbf{x}^{(m)})]$  is invertible.

The following assumption requires that the support of the distribution  $\mathbb{P}_{\mathbf{W}}$  is sufficiently rich.

**Assumption 3.4** (Sufficient task variability). There exists  $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(m)} \in \mathcal{W}$  and indices  $i_1, \dots, i_m \in [k]$  such that the rows  $\mathbf{W}_{i_1, \cdot}^{(1)}, \dots, \mathbf{W}_{i_m, \cdot}^{(m)}$  are linearly independent.

Under Assumptions 3.2 to 3.4, the representation  $\mathbf{f}_\theta$  is identifiable up to linear equivalence (see Theorem B.4 in Appendix B). Similar results were shown by Roeder et al. (2021); Ahuja et al. (2022c). The next assumptions will guarantee disentanglement.

In order to formalize the intuitive idea that most tasks do not require all features, we will denote by  $S^{(t)}$  the support of the matrix  $\mathbf{W}^{(t)}$ , i.e.,

$$S^{(t)} := \{j \in [m] \mid \mathbf{W}_{:,j}^{(t)} \neq \mathbf{0}\}.$$

In other words,  $S^{(t)}$  is the set of features which are useful to predict  $y$  in the  $t$ -th task; note that it is unknown to the learner. For our analysis, we decompose  $\mathbb{P}_{\mathbf{W}}$  as

$$\mathbb{P}_{\mathbf{W}} = \sum_{S \in \mathcal{P}([m])} p(S) \mathbb{P}_{\mathbf{W}|S}, \quad (5)$$

where  $\mathcal{P}([m])$  is the collection of all subsets of  $[m]$ ,  $p(S)$  is the probability that the support of  $\mathbf{W}$  is  $S$  and  $\mathbb{P}_{\mathbf{W}|S}$  is the conditional distribution of  $\mathbf{W}$  given that its support is  $S$ . Let  $\mathcal{S}$  be the support of the distribution  $p(S)$ , i.e.,  $\mathcal{S} := \{S \in \mathcal{P}([m]) \mid p(S) > 0\}$ . The set  $\mathcal{S}$  will have an important role in Assumption 3.6.

The following assumption requires that  $\mathbb{P}_{\mathbf{W}|S}$  does not concentrate mass on certain proper subspaces.

**Assumption 3.5** (Intra-support sufficient task variability). For all  $S \in \mathcal{S}$  and all  $\mathbf{a} \in \mathbb{R}^{|S|} \setminus \{0\}$ ,

$$\mathbb{P}_{\mathbf{W}|S} \{\mathbf{W} \in \mathbb{R}^{k \times m} \mid \mathbf{W}_{:,S} \mathbf{a} = \mathbf{0}\} = 0.$$

We illustrate the above assumption in the simpler case where  $k = 1$ . For instance, Assumption 3.5 holds when the distribution of  $\mathbf{W}_{1,S} \mid S$  has a density w.r.t. the Lebesgue measure on  $\mathbb{R}^{|S|}$ , which is true for example when  $\mathbf{W}_{1,S} \mid$

$S \sim \mathcal{N}(\mathbf{0}, \Sigma)$  and the covariance matrix  $\Sigma$  is full rank (red distribution in Fig. 2). However, if  $\Sigma$  is not full rank, the probability distribution of  $\mathbf{W}_{1,S} \mid S$  concentrates its mass on a proper linear subspace  $V \subsetneq \mathbb{R}^{|S|}$ , which violates Assumption 3.5 (blue distribution in Fig. 2). Another important counter-example is when  $\mathbb{P}_{\mathbf{W}|S}$  concentrates some of its mass on a point  $\mathbf{W}^{(0)}$ , i.e.,  $\mathbb{P}_{\mathbf{W}|S}\{\mathbf{W}^{(0)}\} > 0$  (orange distribution in Fig. 2). We provide a concrete numerical example of what can go wrong when the support of the  $\mathbb{P}_{\mathbf{W}|S}$  is finite in Appendix B.4. Interestingly, there are distributions over  $\mathbf{W}_{1,S} \mid S$  that do not have a density w.r.t. the Lebesgue measure, but still satisfy Assumption 3.5. This is the case, e.g., when  $\mathbf{W}_{1,S} \mid S$  puts uniform mass over a  $(|S| - 1)$ -dimensional sphere embedded in  $\mathbb{R}^{|S|}$  and centered at zero. See Appendix B.6 for a justification.

The following assumption requires that the support  $\mathcal{S}$  of  $p(S)$  is “rich enough”.

**Assumption 3.6** (Sufficient variability of the task supports). For all  $j \in [m]$ ,

$$\bigcup_{S \in \mathcal{S} | j \notin S} S = [m] \setminus \{j\}.$$

Intuitively, Assumption 3.6 requires that, for every feature  $j$ , one can find a set of tasks such that their supports cover all features except  $j$  itself. Fig. 3 shows an example of  $\mathcal{S}$  satisfying Assumption 3.6. Appendix B.5 provides a probabilistic argument showing that Assumption 3.6 holds “in most cases” when the number of supports is very large. That being said, we conjecture that removing this assumption would yield a form of *partial disentanglement* resembling the one developed by Lachapelle & Lacoste-Julien (2022) in which some groups of latent factors would remain entangled.

### 3.4. Tractable bilevel optimization problems for sparse multitask learning

The proposed approach to jointly estimate the representation and the task-specific predictors relies on a bilevel optimization problem (Problem (4)) that is intractable because of the non-convex constraints. To obtain a tractable bi-level optimization problem, the  $L_{2,0}$  constraints are replaced by their convex relaxations in the penalized form, which are also known to promote group sparsity (Argyriou et al., 2008):

$$\begin{aligned} \min_{\hat{\theta}} \quad & -\frac{1}{Tn} \sum_{t=1}^T \sum_{(\mathbf{x}, y) \in \mathcal{D}_t} \log p(y; \hat{\mathbf{W}}^{(t)} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \\ \text{s.t.} \quad & \hat{\mathbf{W}}^{(t)} \in \arg \min_{\tilde{\mathbf{W}}} \frac{1}{n} \sum_{(\mathbf{x}, y) \in \mathcal{D}_t} -\log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \\ & \quad \quad \quad + \lambda_t \|\tilde{\mathbf{W}}\|_{2,1}, \end{aligned} \quad (6)$$

where the constraint holds for all  $t \in [T]$ . Following Benio (2000); Pedregosa (2016), one can compute the (hyper)gradient of the outer function using implicit differentiation, even if the inner optimization problem is non-smooth (Bertrand et al., 2020; Bolte et al., 2021; Malézieux et al., 2022; Bolte et al., 2022). Once the hypergradient is computed, one can optimize Problem (6) with usual first-order methods (Wright & Nocedal, 1999).

Note that the quantity  $\hat{\mathbf{W}}^{(t)} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$  is invariant to simultaneous rescaling of  $\hat{\mathbf{W}}^{(t)}$  by a scalar and of  $\mathbf{f}_{\hat{\theta}}(\mathbf{x})$  by its inverse. Thus, without constraints on  $\mathbf{f}_{\hat{\theta}}(\mathbf{x})$ ,  $\|\hat{\mathbf{W}}^{(t)}\|_{2,1}$  can be made arbitrarily small. This issue is similar to the one faced in sparse dictionary learning (Kreutz-Delgado et al., 2003; Mairal et al., 2008; 2009; 2011), where unit-norm constraints are usually imposed on dictionary columns. In our case, since  $\mathbf{f}_{\hat{\theta}}$  is parametrized by a neural network, we suggest applying batch or layer normalization (Ioffe & Szegedy, 2015; Ba et al., 2016) to control the norm of  $\mathbf{f}_{\hat{\theta}}(\mathbf{x})$ . Since the number of relevant features might be task-dependent, Problem (6) has one regularization hyperparameter  $\lambda_t$  per task. However, in practice, we select  $\lambda_t := \lambda$  for all  $t \in [T]$  to limit the number of hyperparameters. We also use an adaptive scheme to have  $\lambda$  in a reasonable range throughout training, which we explain in Appendix D.2.3.

Appendix B.3 introduces a similar relaxation of Theorem B.6 (mentioned in Section 3.2) in which the sparsity penalty appears in the outer problem instead of the inner problem. Appendix D.2.5 presents empirical results showing this alternative approach yields very similar results.

**Link with meta-learning.** The bi-level formulation Problem (6) is closely related to *metric-based meta-learning* methods (Snell et al., 2017; Bertinetto et al., 2019), where a shared representation  $\mathbf{f}_{\hat{\theta}}$  is learned across all tasks via simple task-specific predictors, such as linear classifiers. In the general meta-learning setting (Finn et al., 2017), one is given a large number of training datasets  $(\mathcal{D}_t^{\text{train}})_{1 \leq t \leq T}$ , which usually only contain a small number of samples  $n$ . As opposed to the multi-task setting (i.e., unlike in Section 3.1), one is also given separate *test datasets*  $(\mathcal{D}_t^{\text{test}})_{1 \leq t \leq T}$  of  $n'$  samples for each task  $t$ , to evaluate how well the learned model generalizes to new test samples. In meta-learning, the goal is to *learn a learning procedure* that will generalize well on new unseen tasks.

Formally, metric-based meta-learning can be formulated as

$$\begin{aligned} \min_{\hat{\theta}} \quad & \frac{1}{Tn'} \sum_{t=1}^T \sum_{(\mathbf{x}, y) \in \mathcal{D}_t^{\text{test}}} \mathcal{L}_{\text{out}}(\hat{\mathbf{W}}_{\hat{\theta}}^{(t)}; \mathbf{f}_{\hat{\theta}}(\mathbf{x}), y) \\ \text{s.t.} \quad & \hat{\mathbf{W}}_{\hat{\theta}}^{(t)} \in \arg \min_{\tilde{\mathbf{W}}} \frac{1}{n} \sum_{(\mathbf{x}, y) \in \mathcal{D}_t^{\text{train}}} \mathcal{L}_{\text{in}}(\tilde{\mathbf{W}}; \mathbf{f}_{\hat{\theta}}(\mathbf{x}), y). \end{aligned} \quad (7)$$

The main difference between [Problem \(6\)](#) and [Problem \(7\)](#) is that, in the latter, the inner and outer loss functions  $\mathcal{L}_{\text{in}}$  and  $\mathcal{L}_{\text{out}}$  are not evaluated on the same dataset. [Section 5.2](#) shows experiments with a meta-learning variant of [Problem \(6\)](#) based on group Lasso multiclass SVM predictors.

## 4. Related Work

**Disentanglement.** Since the work of [Bengio et al. \(2013\)](#), many methods have been proposed to learn disentangled representations based on various heuristics ([Higgins et al., 2017](#); [Chen et al., 2018](#); [Kim & Mnih, 2018](#); [Kumar et al., 2018](#); [Bouchacourt et al., 2018](#)). Following the work of [Locatello et al. \(2019\)](#), which highlighted the lack of identifiability in modern deep generative models, many works have proposed more or less weak forms of supervision motivated by identifiability analyses ([Locatello et al., 2020a](#); [Klindt et al., 2021](#); [Von Kügelgen et al., 2021](#); [Ahuja et al., 2022a;c](#); [Zheng et al., 2022](#)). A similar line of work have adopted the causal representation learning perspective ([Lachapelle et al., 2022](#); [Lachapelle & Lacoste-Julien, 2022](#); [Lippe et al., 2022b;a](#); [Ahuja et al., 2022b](#); [Yao et al., 2022](#); [Brehmer et al., 2022](#)).

The problem of identifiability was well known among the *independent component analysis* (ICA) community ([Hyvärinen et al., 2001](#); [Hyvärinen & Pajunen, 1999](#)) which came up with solutions for general nonlinear mixing functions by leveraging auxiliary information ([Hyvärinen & Morioka, 2016](#); [2017](#); [Hyvärinen et al., 2019](#); [Khemakhem et al., 2020a;b](#)). Another approach is to consider restricted hypothesis classes of mixing functions ([Taleb & Jutten, 1999](#); [Gresele et al., 2021](#); [Zheng et al., 2022](#); [Moran et al., 2022](#)). [Locatello et al. \(2020b\)](#) proposed a semi-supervised learning approach to disentangle in cases where a few samples are labelled with the values of the factors of variations themselves. This is different from our approach as the labels that we consider can be sampled from some  $p(y; \mathbf{W} \mathbf{f}_{\hat{\theta}}(\mathbf{x}))$ , which is more general. [Ahuja et al. \(2022c\)](#) consider a setting similar to ours, but they rely on the independence and non-gaussianity of the latent factors for disentanglement using linear ICA. See the end of [Section 3.2](#) for further discussions on how our theory distinguishes itself from most methods cited above.

**Multi-task, transfer & invariant learning.** While the statistical advantages of multi-task representation learning are well understood ([Lounici et al., 2011a;b](#); [Maurer et al., 2016](#)), the theoretical benefits of disentanglement for transfer learning are not clearly established (apart from [Zhang et al. 2022](#)). Some works have investigated this question empirically and obtained both positive ([van Steenkiste et al., 2019](#); [Miladinovi et al., 2019](#); [Dittadi et al., 2021](#)) and negative results ([Locatello et al., 2019](#); [Montero et al., 2021](#)). Invariant risk minimization ([Arjovskiy et al.,](#)

[2020](#); [Ahuja et al., 2020](#); [Krueger et al., 2021](#); [Lu et al., 2021](#)) aims at learning a representation that elicits a single predictor that is optimal for all tasks. This differs from our approach which learns one predictor per task.

**Dictionary learning and sparse coding.** We contrast our approach, which jointly learns a *dense representation* and sparse task-specific predictors ([Problem \(6\)](#)), with the line of work which consists in learning *sparse representations* ([Chen et al., 1998](#); [Gribonval & Lesage, 2006](#)). For instance, sparse dictionary learning ([Mairal et al., 2009](#); [2011](#); [Maurer et al., 2013](#)) is an unsupervised technique that aims at learning a dictionary of *atoms* used to reconstruct inputs via sparse linear combinations of its elements. The representation of a single input consists of the coefficients of the linear combination of atoms that minimizes a sparsity-regularized reconstruction loss. In the case of supervised dictionary learning ([Mairal et al., 2008](#)), an additional (potentially expressive) classifier is learned on top of that representation. This large literature has led to a wide variety of estimators: for instance, [Mairal et al. \(2008, Eq. 4\)](#), which minimizes the sum of the classification error and the approximation error of the code, or [Mairal et al. \(2011\)](#), introducing bi-level formulations. While sharing similar optimization challenges, our method is conceptually different and computes the representation of a single input  $\mathbf{x}$  by evaluating the learned function  $\mathbf{f}_{\hat{\theta}}$ .

## 5. Experiments

We present experiments on disentanglement and few-shot learning. Our implementation relies on `jax` and `jaxopt` ([Bradbury et al., 2018](#); [Blondel et al., 2022](#)) and is available here: <https://github.com/tristandeleu/synergies-disentanglement-sparsity>.

### 5.1. Disentanglement in 3D Shapes

We now illustrate [Theorem 3.1](#) by applying [Problem \(6\)](#) to tasks generated using the 3D Shapes dataset ([Burgess & Kim, 2018](#)).

*Data generation.* For all tasks  $t$ , the labelled dataset  $\mathcal{D}_t = \{(\mathbf{x}^{(t,i)}, y^{(t,i)})\}_{i=1}^n$  is generated by first sampling the ground-truth latent variables  $\mathbf{z}^{(t,i)}$  i.i.d. according to some distribution  $p(\mathbf{z})$ , while the corresponding input is obtained doing  $\mathbf{x}^{(t,i)} := \mathbf{f}_{\theta}^{-1}(\mathbf{z}^{(t,i)})$  ( $\mathbf{f}_{\theta}$  is invertible in 3D Shapes). Then, a sparse weight vector  $\mathbf{w}^{(t)}$  is sampled randomly to compute the labels of each example as  $y^{(t,i)} := \mathbf{w}^{(t)} \cdot \mathbf{z}^{(t,i)} + \epsilon^{(t,i)}$ , where  $\epsilon^{(t,i)}$  is independent Gaussian noise. [Fig. 4](#) explores various choices of  $p(\mathbf{z})$  by varying the level of correlation between the latent variables and by varying the level of noise on the ground-truth latents. See [Appendix D.2](#) for more details about the data generating process and [Fig. 7](#) to visualize various  $p(\mathbf{z})$ .




Figure 4. Disentanglement performance (MCC) for all three methods considered as a function of the regularization parameter (left and middle). Varying level of correlation between latents (top) and noise on the latents (bottom). The right columns show performances of the best hyperparameter for different values of correlation and noise. We explain what is  $\lambda_{\max}$  in Appendix D.2.3.

*Algorithms.* In this setting where  $p(y; \boldsymbol{\eta})$  is a Gaussian with fixed variance, the inner problem of Problem (6) amounts to Lasso regression, we thus refer to this approach as inner-Lasso. We also evaluate a simple variation of Problem (6) in which the  $L_1$  norm is replaced by an  $L_2$  norm and refer to it as inner-Ridge. In addition, we evaluate the representation obtained by performing linear ICA (Comon, 1992) on the representation learned by inner-Ridge: the case  $\lambda = 0$  corresponds to the approach of Ahuja et al. (2022c).

*Discussion.* Fig. 4 reports disentanglement performances of the three methods, as measured by the *mean correlation coefficient*, or MCC (Hyvärinen & Morioka, 2016; Khe-makhem et al., 2020a) (Appendix D.2). In all settings, inner-Lasso obtains high MCC for some values of  $\lambda$ , being on par or surpassing the baselines. As the theory suggests, it is robust to high levels of correlations between the latents, as opposed to inner-Ridge with ICA which is very much affected by strong correlations (since ICA assumes independence). We can also see how additional noise on the latent variables hurts inner-Ridge with ICA while leaving inner-Lasso unaffected. Fig. 6 in Appendix D.2 shows that all methods find a representation which is linearly equivalent to the ground-truth representation, except for very large values of  $\lambda$ . Appendix D.2.4 studies empirically to what extent inner-Lasso is robust to violations of Assumption 3.6, Appendix D.2.6 presents a visual evaluation of disentanglement and Appendix D.2.7 reports the DCI metric (Eastwood & Williams, 2018) on the same experiments. We did not explore hyperparameter selection in this work, which is a difficult problem for disentanglement because a goodness-of-fit score evaluated on a held-out dataset will not be informative because of the lack of identifiability. Nevertheless, one can use heuristics such as the *unsupervised disentanglement ranking* score proposed by Duan et al. (2020).

## 5.2. Sparse task-specific predictors in few-shot learning

Despite the lack of ground-truth latent factors in standard few-shot learning benchmarks, we also evaluate sparse meta-learning objectives on the *miniImageNet* dataset (Vinyals et al., 2016). The purpose of this experiment is to show that the sparse formulation of standard metric-based meta-learning techniques reaches similar performance while using a fraction of the features (Fig. 5, right).

Inspired by Lee et al. (2019), where the task-specific classifiers are multiclass support-vector machines (SVMs, Crammer & Singer 2001), we propose to use group Lasso penalized multiclass SVMs, to introduce sparsity in the classifiers. Using the notation of Problem (7), we choose

$$\mathcal{L}_{\text{in}}(\mathbf{W}; f_{\boldsymbol{\theta}}(\mathbf{x}_i), \mathbf{y}_i) = \max_{l \in [k]} ((\mathbf{W}_{\mathbf{y}_i} - \mathbf{W}_l) \cdot f_{\boldsymbol{\theta}}(\mathbf{x}_i) - \mathbf{Y}_{il}) + \lambda_1 \|\mathbf{W}\|_{2,1} + \frac{\lambda_2}{2} \|\mathbf{W}\|^2, \quad (8)$$

$$\mathcal{L}_{\text{out}}(\mathbf{W}; f_{\boldsymbol{\theta}}(\mathbf{x}_i), \mathbf{y}_i) = \text{CE}(\mathbf{W} f_{\boldsymbol{\theta}}(\mathbf{x}_i), \mathbf{Y}_i), \quad (9)$$

with  $\mathbf{Y} \in \mathbb{R}^{n \times k}$  the one-hot encoding of  $\mathbf{y} \in \mathbb{R}^n$  and CE the cross-entropy. The difference with Lee et al. (2019) is the sparsity-promoting term  $\|\mathbf{W}\|_{2,1}$ , which makes the bi-level optimization problem harder to solve. That is why we propose solving the dual (Boyd et al., 2004, Chap. 5) of this inner optimization problem, which writes

$$\min_{\boldsymbol{\Lambda} \in \mathbb{R}^{n \times k}} \frac{1}{\lambda_2} \sum_{j=1}^m \|\text{BST}((\mathbf{Y} - \boldsymbol{\Lambda})^\top \mathbf{F}_{:j}, \lambda_1)\|^2 + \langle \mathbf{Y}, \boldsymbol{\Lambda} \rangle$$

$$\text{s.t. } \forall i, l, \in [n] \times [k], \sum_{l'=1}^k \boldsymbol{\Lambda}_{il'} = 1 \text{ and } \boldsymbol{\Lambda}_{il} \geq 0, \quad (10)$$

with  $\text{BST} : (\mathbf{a}, \tau) \mapsto (1 - \tau / \|\mathbf{a}\|)_+ \mathbf{a}$  is the block soft-thresholding operator,  $\mathbf{F} \in \mathbb{R}^{n \times m}$  the concatenation of



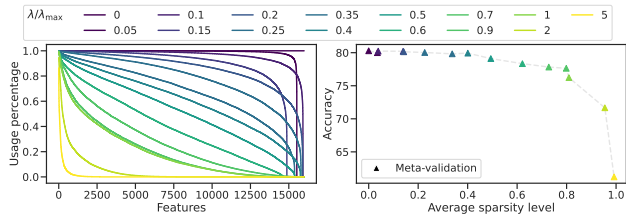


Figure 5. *Left.* Effect of sparsity on the percentage of tasks using specific features, with our meta-learning objective, on *miniImageNet*. *Right.* The meta-validation accuracy of the meta-learning algorithm against the average level of sparsity in the task-specific predictor, for different values of  $\lambda$ .

$\{\mathbf{f}_{\hat{\theta}}(x)\}_{(x,y) \in \mathcal{D}^{\text{train}}}$ . In addition, the primal-dual link writes,  $\forall j \in [m]$ ,  $\mathbf{W}_{:j} = \text{BST}((\mathbf{Y} - \mathbf{\Lambda})^\top \mathbf{F}_{:j}, \lambda_1) / \lambda_2$ . The derivation of the dual can be found in Appendix C.1, Solving this kind of problem in the dual is standard in the SVM literature: it has been proven to be computationally advantageous (Hsieh et al., 2008) when the number of features  $m$  is significantly larger than the number of samples  $n$  (here  $m = 1.6 \times 10^4$  and  $n \leq 25$ ). Details on how to solve and differentiate through Problem (10) are in Appendix D.3.

*Discussion.* In Fig. 5 (right), we observe that the accuracy of the sparse meta-learning method on novel (meta-validation) tasks is similar to the dense counterpart ( $\lambda = 0$ ), while using only a few of the features available (around 30% of sparsity, with no impact on the performance). Naturally, the performance starts to drop as the sparsity level increases though, albeit being still competitive. We also report in Fig. 5 (left) how frequently each feature in the learned representation is used by the task-specific predictors on meta-validation tasks (sorted by usage, for each  $\lambda$ ). The gradual decrease in usage suggests that the features are reused in different contexts, across different tasks.

## 6. Conclusion

In this work, we investigated the synergies between sparsity, disentanglement and generalization. We showed that when the downstream task can be solved using only a fraction of the factors of variations, disentangled representations combined with sparse task-specific predictors can improve generalization (Section 2). Our novel identifiability result (Theorem 3.1) sheds light on how, in a multi-task setting, sparsity regularization on the task-specific predictors can induce disentanglement. This led to a practical bi-level optimization problem that was shown to yield disentangled representations on regression tasks based on the 3D Shapes dataset. Finally, we explored the connection between this bi-level formulation and meta-learning, and we showed how sparse task-specific predictors may achieve similar performance on unseen tasks with only a fraction of the features. Future work could explore identifiability in a

more general setting where the task-specific predictors are potentially nonlinear, which should be applicable to more problems.

## Acknowledgements

This research was partially supported by the Canada CIFAR AI Chair Program, by an IVADO excellence PhD scholarship and by a Google Focused Research award. The experiments were in part enabled by computational resources provided by Calcul Quebec and Compute Canada. Simon Lacoste-Julien is a CIFAR Associate Fellow in the Learning in Machines & Brains program. Sébastien Lachapelle and Quentin Bertrand would like to thank Samsung Electronics Co., Ltd. for funding this research. The authors would like to thank David Berger for insightful discussions in the early stage of this project.

## References

- Ahuja, K., Shanmugam, K., Varshney, K. R., and Dhurandhar, A. Invariant risk minimization games. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Ahuja, K., Hartford, J., and Bengio, Y. Properties from mechanisms: an equivariance perspective on identifiable representation learning. In *International Conference on Learning Representations*, 2022a.
- Ahuja, K., Hartford, J., and Bengio, Y. Weakly supervised representation learning with sparse perturbations, 2022b.
- Ahuja, K., Mahajan, D., Syrgkanis, V., and Mitliagkas, I. Towards efficient representation identification in supervised learning. In *First Conference on Causal Learning and Reasoning*, 2022c.
- Argyriou, A., Evgeniou, T., and Pontil, M. Convex multi-task feature learning. *Machine learning*, 73(3):243–272, 2008.
- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization, 2020.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bengio, Y. Gradient-based optimization of hyperparameters. *Neural computation*, 12(8):1889–1900, 2000.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 2013.
- Bertinetto, L., Henriques, J. F., Torr, P. H., and Vedaldi, A. Meta-learning with differentiable closed-form solvers. 2019.

- Bertrand, Q., Klopfenstein, Q., Blondel, M., Vaiter, S., Gramfort, A., and Salmon, J. Implicit differentiation of lasso-type models for hyperparameter optimization. In *International Conference on Machine Learning*, pp. 810–821. PMLR, 2020.
- Bertrand, Q., Klopfenstein, Q., Massias, M., Blondel, M., Vaiter, S., Gramfort, A., and Salmon, J. Implicit differentiation for fast hyperparameter selection in non-smooth convex learning. *JMLR*, 2022.
- Bickel, P. J., Ritov, Y., and Tsybakov, A. B. Simultaneous analysis of lasso and Dantzig selector. *The Annals of statistics*, 37(4):1705–1732, 2009.
- Blondel, M., Berthet, Q., Cuturi, M., Frostig, R., Hoyer, S., Llinares-López, F., Pedregosa, F., and Vert, J.-P. Efficient and modular implicit differentiation. *NeurIPS*, 2022.
- Bolte, J., Le, T., E., Pauwels, and Silveti-Falls, T. Nonsmooth implicit differentiation for machine-learning and optimization. *Advances in neural information processing systems*, 34:13537–13549, 2021.
- Bolte, J., Pauwels, E., and Vaiter, S. Automatic differentiation of nonsmooth iterative algorithms. *NeurIPS*, 2022.
- Bouchacourt, D., Tomioka, R., and Nowozin, S. Multi-level variational autoencoder: Learning disentangled representations from grouped observations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- Boyd, S. P., , and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Brehmer, J., De Haan, P., Lippe, P., and Cohen, T. Weakly supervised causal representation learning. In *Advances in Neural Information Processing Systems*, 2022.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, 2020a.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33: 1877–1901, 2020b.
- Burgess, C. and Kim, H. 3d shapes dataset. <https://github.com/deepmind/3dshapes-dataset/>, 2018.
- Casella, G. and Berger, R. *Statistical Inference*. Duxbury Resource Center, 2001.
- Chen, R. T. Q., Li, X., G., R., and Duvenaud, D. Isolating sources of disentanglement in vaes. In *Advances in Neural Information Processing Systems*, 2018.
- Chen, S. S., Donoho, D. L., and Saunders, M. A. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 1998.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pp. 1597–1607. PMLR, 2020.
- Comon, P. Independent component analysis. *Higher-Order Statistics*, 1992.
- Crammer, K. and Singer, Y. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of machine learning research*, 2(Dec):265–292, 2001.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Dittadi, A., Träuble, F., Locatello, F., Wuthrich, M., Agrawal, V., Winther, O., Bauer, S., and Schölkopf, B. On the transfer of disentangled representations in realistic settings. In *International Conference on Learning Representations*, 2021.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations*, 2021.
- Duan, S., Matthey, L., Saraiva, A., Watters, N., Burgess, C., Lerchner, A., and Higgins, I. Unsupervised model selection for variational disentangled representation learning. In *International Conference on Learning Representations*, 2020.
- Eastwood, C. and Williams, C. K. A framework for the quantitative evaluation of disentangled representations. In *International Conference on Learning Representations*, 2018.

- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.
- Goyal, A. and Bengio, Y. Inductive biases for deep learning of higher-level cognition. *Proc. R. Soc. A* 478: 20210068, 2022.
- Gresele, L., Kügelgen, J. V., Stimper, V., Schölkopf, B., and Besserve, M. Independent mechanism analysis, a new concept? In *Advances in Neural Information Processing Systems*, 2021.
- Gribonval, R. and Lesage, S. A survey of sparse component analysis for blind source separation: principles, perspectives, and new challenges. In *ESANN'06 proceedings - 14th European Symposium on Artificial Neural Networks*, 2006.
- Higgins, I., Matthey, L., Pal, A., Burgess, C. P., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017.
- Hoerl, A. E. and Kennard, R. W. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- Hospedales, T., Antoniou, A., Micaelli, P., and Storkey, A. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5149–5169, 2021.
- Hsieh, C.-J., Chang, K.-W., Lin, C.-J., Keerthi, S. S., and Sundararajan, S. A dual coordinate descent method for large-scale linear svm. In *Proceedings of the 25th international conference on Machine learning*, pp. 408–415, 2008.
- Hyvärinen, A. and Morioka, H. Unsupervised feature extraction by time-contrastive learning and nonlinear ica. In *Advances in Neural Information Processing Systems*, 2016.
- Hyvärinen, A. and Morioka, H. Nonlinear ICA of Temporally Dependent Stationary Sources. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017.
- Hyvärinen, A. and Pajunen, P. Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 1999.
- Hyvärinen, A., Karhunen, J., and Oja, E. *Independent Component Analysis*. Wiley, 2001.
- Hyvärinen, A., Sasaki, H., and Turner, R. E. Nonlinear ica using auxiliary variables and generalized contrastive learning. In *AISTATS*. PMLR, 2019.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning*, 2015.
- Khemakhem, I., Kingma, D., Monti, R., and Hyvärinen, A. Variational autoencoders and nonlinear ica: A unifying framework. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, 2020a.
- Khemakhem, I., Monti, R., Kingma, D., and Hyvärinen, A. Ice-beem: Identifiable conditional energy-based deep models based on nonlinear ica. In *Advances in Neural Information Processing Systems*, 2020b.
- Kim, H. and Mnih, A. Disentangling by factorising. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Klindt, D. A., Schott, L., Sharma, Y., Ustyuzhaninov, I., Brendel, W., Bethge, M., and Paiton, D. M. Towards nonlinear disentanglement in natural data with temporal sparse coding. In *9th International Conference on Learning Representations*, 2021.
- Kreutz-Delgado, K., Murray, J. F., Rao, B. D., Engan, K., Lee, T.-W., and Sejnowski, T. J. Dictionary learning algorithms for sparse representation. *Neural computation*, 15(2):349–396, 2003.
- Krueger, D., Caballero, E., Jacobsen, J.-H., Zhang, A., Binas, J., Priol, R. L., Zhang, D., and Courville, A. Out-of-distribution generalization via risk extrapolation ( $\{r\}x$ ), 2021.
- Kumar, A., Sattigeri, P., and Balakrishnan, A. Variational inference of disentangled latent concepts from unlabeled observations. In *International Conference on Learning Representations*, 2018.
- Lachapelle, S. and Lacoste-Julien, S. Partial disentanglement via mechanism sparsity. In *UAI 2022 Workshop on Causal Representation Learning*, 2022.
- Lachapelle, S., Rodriguez Lopez, P., Sharma, Y., Everett, K. E., Le Priol, R., Lacoste, A., and Lacoste-Julien, S. Disentanglement via mechanism sparsity regularization: A new principle for nonlinear ICA. In *First Conference on Causal Learning and Reasoning*, 2022.
- Lee, K., S.Maji, Ravichandran, A., and Soatto, S. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10657–10665, 2019.

- Lippe, P., Magliacane, S., Löwe, S., Asano, Y. M., Cohen, T., and Gavves, E. iCITRIS: Causal representation learning for instantaneous temporal effects. In *UAI 2022 Workshop on Causal Representation Learning*, 2022a.
- Lippe, P., Magliacane, S., Löwe, S., Asano, Y. M., Cohen, T., and Gavves, E. CITRIS: Causal identifiability from temporal intervened sequences, 2022b.
- 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 *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- Locatello, F., Poole, B., Raetsch, G., Schölkopf, B., Bachem, O., and Tschannen, M. Weakly-supervised disentanglement without compromises. In *Proceedings of the 37th International Conference on Machine Learning*, 2020a.
- Locatello, F., Tschannen, M., Bauer, S., Räscht, G., Schölkopf, B., and Bachem, O. Disentangling factors of variations using few labels. In *International Conference on Learning Representations*, 2020b. URL <https://openreview.net/forum?id=SygagpEKwB>.
- Lounici, K., Pontil, M., and Tsybakov, A. B. Oracle inequalities and optimal inference under group sparsity. *The Annals of statistics*, 2011a.
- Lounici, K., Pontil, M., Van De Geer, S., and Tsybakov, A. B. Oracle inequalities and optimal inference under group sparsity. *The annals of statistics*, 2011b.
- Lu, C., Wu, Y., Hernández-Lobato, J. M., and Schölkopf, B. Nonlinear invariant risk minimization: A causal approach, 2021.
- Mairal, J., Ponce, J., Sapiro, G., Zisserman, A., and Bach, F. Supervised dictionary learning. *Advances in neural information processing systems*, 21, 2008.
- Mairal, J., Bach, F., Ponce, J., and Sapiro, G. Online dictionary learning for sparse coding. In *Proceedings of the 26th annual international conference on machine learning*, pp. 689–696, 2009.
- Mairal, J., Bach, F., and Ponce, J. Task-driven dictionary learning. *IEEE transactions on pattern analysis and machine intelligence*, 34(4):791–804, 2011.
- Malézieux, B., Moreau, T., and Kowalski, M. Dictionary and prior learning with unrolled algorithms for unsupervised inverse problems. *ICLR*, 2022.
- Marcus, G., Davis, E., and Aaronson, S. A very preliminary analysis of dall-e 2. *arXiv preprint arXiv:2204.13807*, 2022.
- Maurer, A., Pontil, M., and Romera-Paredes, B. Sparse coding for multitask and transfer learning. *ICML'13*, 2013.
- Maurer, A., Pontil, M., and Romera-Paredes, B. The benefit of multitask representation learning. *J. Mach. Learn. Res.*, 2016.
- Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., and Khudanpur, S. Recurrent neural network based language model. *ISCA*, 2010.
- Miladinovi, D., Gondal, M. W., Schölkopf, B., Buhmann, J. M., and Bauer, S. Disentangled state space representations. *arXiv preprint arXiv:1906.03255*, 2019.
- Mohri, M., Rostamizadeh, A., and Talwalkar, A. MIT Press, 2018.
- Montero, M. L., Ludwig, C. J., Costa, R. P., Malhotra, G., and Bowers, J. The role of disentanglement in generalisation. In *International Conference on Learning Representations*, 2021.
- Moran, G. E., Sridhar, D., Wang, Y., and Blei, D. Identifiable deep generative models via sparse decoding. *Transactions on Machine Learning Research*, 2022.
- Oord, A., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *Advances in Neural Information Processing Systems*, 2018.
- Pearl, J. The seven tools of causal inference, with reflections on machine learning. *Commun. ACM*, 2019.
- Pedregosa, F. Hyperparameter optimization with approximate gradient. In *International conference on machine learning*, pp. 737–746. PMLR, 2016.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pp. 8748–8763. PMLR, 2021.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- Richtárik, P. and Takáč, M. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Mathematical Programming*, 144(1):1–38, 2014.
- Roeder, G., Metz, L., and Kingma, D. P. On linear identifiability of learned representations. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.

- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., and Bengio, Y. Toward causal representation learning. *Proceedings of the IEEE - Advances in Machine Learning and Deep Neural Networks*, 2021.
- Schölkopf, B. Causality for machine learning, 2019.
- Snell, J., Swersky, K., and Zemel, R. Prototypical networks for few-shot learning. *Advances in Neural Information Processing Systems*, 30, 2017.
- Taleb, A. and Jutten, C. Source separation in post-nonlinear mixtures. *IEEE Transactions on Signal Processing*, 1999.
- Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Tseng, P. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of optimization theory and applications*, 109(3):475–494, 2001.
- van Steenkiste, S., Locatello, F., Schmidhuber, J., and Bachem, O. Are disentangled representations helpful for abstract visual reasoning? In *Advances in Neural Information Processing Systems*, 2019.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016.
- 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. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- Wainwright, M. J. and Jordan, M. I. Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.*, 2008.
- Wortsman, M., Ilharco, G., Kim, J. W., Li, M., Kornblith, S., Roelofs, R., Lopes, R. G., Hajishirzi, H., Farhadi, A., Namkoong, H., and Schmidt, L. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7959–7971, June 2022.
- Wright, S. and Nocedal, J. Numerical optimization. *Springer Science*, 35(67-68):7, 1999.
- Yao, W., Sun, Y., Ho, A., Sun, C., and Zhang, K. Learning temporally causal latent processes from general temporal data. In *International Conference on Learning Representations*, 2022.
- Zhang, H., Zhang, Y.-F., Liu, W., Weller, A., Schölkopf, B., and Xing, E. Towards principled disentanglement for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- Zheng, Y., Ng, I., and Zhang, K. On the identifiability of nonlinear ICA: Sparsity and beyond. In *Advances in Neural Information Processing Systems*, 2022.

# Appendix

Table 1. Table of Notations.

<u>Norms &amp; pseudonorms</u>	
$\ \cdot\ $	Euclidean norm on vectors and Frobenius norm on matrices
$\ \mathbf{A}\ _{2,1}$	$:= \sum_{j=1}^m \ \mathbf{A}_{:,j}\ $
$\ \mathbf{A}\ _{2,0}$	$:= \sum_{j=1}^m \mathbb{1}_{\ \mathbf{A}_{:,j}\  \neq 0}$ , where $\mathbb{1}$ is the indicator function.
<u>Data</u>	
$\mathbf{x} \in \mathbb{R}^d$	Observations
$\mathcal{X} \subset \mathbb{R}^d$	Support of observations
$y \in \mathbb{R}$	Target
$\mathcal{Y} \subset \mathbb{R}$	Support of targets
<u>Learned/ground-truth model</u>	
$\mathbf{W} \in \mathbb{R}^{k \times m}$	Ground-truth coefficients
$\hat{\mathbf{W}} \in \mathbb{R}^{k \times m}$	Learned coefficients
$\boldsymbol{\theta}$	Ground-truth parameters of the representation
$\hat{\boldsymbol{\theta}}$	Learned parameters of the representation
$f_{\boldsymbol{\theta}} : \mathbb{R}^d \rightarrow \mathbb{R}^m$	Ground-truth representation
$f_{\hat{\boldsymbol{\theta}}} : \mathbb{R}^d \rightarrow \mathbb{R}^m$	Learned representation
$\boldsymbol{\eta} \in \mathbb{R}^k$	Parameter of the distribution $p(y; \boldsymbol{\eta})$
$\mathbb{P}_{\mathbf{W}}$	Distribution over ground-truth coefficient matrices $\mathbf{W}$
$S$	$:= \{j \in [m] \mid \mathbf{W}_{:,j} \neq \mathbf{0}\}$ (support of $\mathbf{W}$ )
$\mathbb{P}_{\mathbf{W} S}$	Conditional distribution of $\mathbf{W}$ given $S$ .
$p(S)$	Ground-truth distribution over possible supports $S$
$S$	Support of the distribution $p(S)$
<u>Optimization</u>	
$W$	Primal variable
$\Lambda$	Dual variable
$h^* : \mathbf{a} \mapsto \sup_{\mathbf{b} \in \mathbb{R}^d} \langle \mathbf{a}, \mathbf{b} \rangle - h(\mathbf{b})$	Fenchel conjugate of the function $h : \mathbb{R}^d \rightarrow \mathbb{R}$
$f \square g : \mathbf{a} \mapsto \min_{\mathbf{b}} f(\mathbf{a} - \mathbf{b}) + g(\mathbf{b})$	inf-convolution of the functions $f$ and $g$
BST : $(\mathbf{a}, \tau)$	$\mapsto (1 - \tau/\ \mathbf{a}\ )_+ \mathbf{a}$ , block soft-thresholding operator

## A. Proofs of Section 2

**Proposition 2.2.** Let  $\hat{\mathbf{W}}_n^{(\hat{\theta})}$  and  $\hat{\mathbf{W}}_n^{(\theta)}$  be the solutions to [Problem \(1\)](#) with the representations  $\mathbf{f}_{\hat{\theta}}$  and  $\mathbf{f}_{\theta}$ , respectively (which we assume are unique). If  $\mathbf{f}_{\hat{\theta}}$  and  $\mathbf{f}_{\theta}$  are linearly equivalent ([Assumption 2.1](#)), then we have,  $\forall \mathbf{x} \in \mathcal{X}$ ,  $\hat{\mathbf{W}}_n^{(\hat{\theta})} \mathbf{f}_{\hat{\theta}}(\mathbf{x}) = \hat{\mathbf{W}}_n^{(\theta)} \mathbf{f}_{\theta}(\mathbf{x})$ .

*Proof.* By definition of  $\hat{\mathbf{W}}^{(\hat{\theta})}$ , we have that, for all  $\hat{\mathbf{W}} \in \mathbb{R}^{k \times m}$ ,

$$\sum_{(\mathbf{x}, y) \in \mathcal{D}} \log p(y; \hat{\mathbf{W}}^{(\hat{\theta})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \geq \sum_{(\mathbf{x}, y) \in \mathcal{D}} \log p(y; \hat{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \quad (11)$$

$$\sum_{(\mathbf{x}, y) \in \mathcal{D}} \log p(y; \hat{\mathbf{W}}^{(\hat{\theta})} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})) \geq \sum_{(\mathbf{x}, y) \in \mathcal{D}} \log p(y; \hat{\mathbf{W}} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})). \quad (12)$$

Because  $\mathbb{R}^{k \times m} \mathbf{L} = \mathbb{R}^{k \times m}$ , we have that, for all  $\hat{\mathbf{W}} \in \mathbb{R}^{k \times m}$ ,

$$\sum_{(\mathbf{x}, y) \in \mathcal{D}} \log p(y; \hat{\mathbf{W}}^{(\hat{\theta})} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})) \geq \sum_{(\mathbf{x}, y) \in \mathcal{D}} \log p(y; \hat{\mathbf{W}} \mathbf{f}_{\theta}(\mathbf{x})), \quad (13)$$

which is to say that  $\hat{\mathbf{W}}^{(\theta)} = \hat{\mathbf{W}}^{(\hat{\theta})} \mathbf{L}$ , or put differently,  $\hat{\mathbf{W}}^{(\hat{\theta})} = \hat{\mathbf{W}}^{(\theta)} \mathbf{L}^{-1}$ . It implies

$$\hat{\mathbf{W}}^{(\hat{\theta})} \mathbf{f}_{\hat{\theta}}(\mathbf{x}) = \hat{\mathbf{W}}^{(\theta)} \mathbf{L}^{-1} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x}) = \hat{\mathbf{W}}^{(\theta)} \mathbf{f}_{\theta}(\mathbf{x}), \quad (14)$$

which is what we wanted to show.  $\square$

**Proposition 2.4.** Let  $\hat{\mathbf{W}}_{\infty}^{(\hat{\theta})}$  be the (assumed unique) solution of the population-based MLE,  $\arg \max_{\tilde{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x}))$ . If [Assumption 2.1](#) (linear equivalence) & [Assumption 2.3](#) (data generating process) hold,  $\hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} = \mathbf{W} \mathbf{L}^{-1}$ .

*Proof.* By definition of  $\hat{\mathbf{W}}_{\infty}^{(\hat{\theta})}$ , we have that, for all  $\tilde{\mathbf{W}} \in \mathbb{R}^{k \times m}$ ,

$$\mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \geq \mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \quad (15)$$

$$\mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})) \geq \mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \tilde{\mathbf{W}} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})). \quad (16)$$

In particular, the inequality holds for  $\tilde{\mathbf{W}} := \mathbf{W} \mathbf{L}^{-1}$ , which yields

$$\mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})) \geq \mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) \quad (17)$$

$$0 \geq \mathbb{E}_{p(\mathbf{x}, y)} \left[ \log p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) - \log p(y; \hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})) \right] \quad (18)$$

$$0 \geq \mathbb{E}_{p(\mathbf{x})} \text{KL}(p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) \parallel p(y; \hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x}))). \quad (19)$$

Since the KL is always non-negative, we have that,

$$\mathbb{E}_{p(\mathbf{x})} \text{KL}(p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) \parallel p(y; \hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x}))) = 0, \quad (20)$$

which in turn implies

$$\mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})) = \mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) \quad (21)$$

$$\mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})) = \mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \mathbf{W} \mathbf{L}^{-1} \mathbf{L} \mathbf{f}_{\theta}(\mathbf{x})) \quad (22)$$

$$\mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) = \mathbb{E}_{p(\mathbf{x}, y)} \log p(y; \mathbf{W} \mathbf{L}^{-1} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \quad (23)$$

$$(24)$$

Since the solution to the population MLE from [Proposition 2.4](#) is assumed to be unique, this equality holds if and only if  $\hat{\mathbf{W}}_{\infty}^{(\hat{\theta})} = \mathbf{W} \mathbf{L}^{-1}$ .  $\square$

## B. Proofs of Section 3

### B.1. Technical Lemmas

The lemmas of this section can be skipped at first read.

The following lemma will be important for proving [Theorem B.5](#). The argument is taken from [Lachapelle et al. \(2022\)](#).

**Lemma B.1** (Sparsity pattern of an invertible matrix contains a permutation). *Let  $\mathbf{L} \in \mathbb{R}^{m \times m}$  be an invertible matrix. Then, there exists a permutation  $\sigma$  such that  $\mathbf{L}_{i,\sigma(i)} \neq 0$  for all  $i$ .*

*Proof.* Since the matrix  $\mathbf{L}$  is invertible, its determinant is non-zero, *i.e.*,

$$\det(\mathbf{L}) := \sum_{\sigma \in \mathfrak{S}_m} \text{sign}(\sigma) \prod_{i=1}^m \mathbf{L}_{i,\sigma(i)} \neq 0, \quad (25)$$

where  $\mathfrak{S}_m$  is the set of  $m$ -permutations. This equation implies that at least one term of the sum is non-zero, meaning there exists  $\sigma \in \mathfrak{S}_m$  such that for all  $i \in [m]$ ,  $\mathbf{L}_{i,\sigma(i)} \neq 0$ .  $\square$

The following technical lemma will help us dealing with almost-everywhere statements and can be safely skipped at a first read. Before presenting it, we recall the formal definition of a support of a distribution.

**Definition B.2.** The support of a Borel measure  $\mu$  over a topological space  $(X, \tau)$  is the set of point  $x \in X$  such that, for all open set  $U \in \tau$  containing  $x$ ,  $\mu(U) > 0$ .

Throughout this work, we assume implicitly that all measures are Borel measures with respect to the standard topology of the space on which they are defined.

**Lemma B.3.** *Assumption 3.4 is equivalent to the following statement: For all  $E_0 \subset \mathbb{R}^{k \times m}$  such that  $\mathbb{P}_{\mathbf{W}}(E_0) = 0$ , there exists  $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(m)} \in \mathcal{W} \setminus E_0$  and indices  $i_1, \dots, i_m \in [k]$  such that the row vectors  $\mathbf{W}_{i_1,:}^{(1)}, \dots, \mathbf{W}_{i_m,:}^{(m)}$  are linearly independent.*

*Proof.* First of all, the " $\Leftarrow$ " direction is trivial since one can simply pick  $E_0 = \emptyset$ .

We now show the " $\Rightarrow$ " direction. First of all, we notice that, since  $\mathbf{W}_{i_1,:}^{(1)}, \dots, \mathbf{W}_{i_m,:}^{(m)}$  are linearly independent, they form a matrix with nonzero determinant, *i.e.*,

$$\det \begin{bmatrix} \mathbf{W}_{i_1,:}^{(1)} \\ \vdots \\ \mathbf{W}_{i_m,:}^{(m)} \end{bmatrix} \neq 0. \quad (26)$$

Define the map  $\eta : (\mathbb{R}^{k \times m})^m \rightarrow \mathbb{R}^{m \times m}$  as

$$\eta(\bar{\mathbf{W}}^{(1)}, \dots, \bar{\mathbf{W}}^{(m)}) := \begin{bmatrix} \bar{\mathbf{W}}_{i_1,:}^{(1)} \\ \vdots \\ \bar{\mathbf{W}}_{i_m,:}^{(m)} \end{bmatrix}, \quad \forall (\bar{\mathbf{W}}^{(1)}, \dots, \bar{\mathbf{W}}^{(m)}) \in (\mathbb{R}^{k \times m})^m, \quad (27)$$

which is continuous. Note that  $\det(\cdot)$  is also a continuous map, hence  $\det \circ \eta$  is continuous as well. Thus, the set  $V := (\det \circ \eta)^{-1}(\mathbb{R} \setminus \{0\})$  is open (since  $\mathbb{R} \setminus \{0\}$  is open). Let  $\mathbb{P}_{\mathbf{W}}^m$  be the product measure over tuples of matrices  $(\bar{\mathbf{W}}^{(1)}, \dots, \bar{\mathbf{W}}^{(m)})$ . Note that its support is  $\mathcal{W}^m$ . Because  $(\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(m)})$  is in the open set  $V$  and in the support of  $\mathbb{P}_{\mathbf{W}}^m$ , we have that

$$0 < \mathbb{P}_{\mathbf{W}}^m(V) \quad (28)$$

$$= \mathbb{P}_{\mathbf{W}}^m(V \cap \mathcal{W}^m) + \mathbb{P}_{\mathbf{W}}^m(V \cap (\mathcal{W}^m)^c) \quad (29)$$

$$\leq \mathbb{P}_{\mathbf{W}}^m(V \cap \mathcal{W}^m) + \mathbb{P}_{\mathbf{W}}^m((\mathcal{W}^m)^c) \quad (30)$$

$$= \mathbb{P}_{\mathbf{W}}^m(V \cap \mathcal{W}^m) \quad (31)$$



Let  $E_0 \subset \mathbb{R}^{k \times m}$  be such that  $\mathbb{P}_{\mathbf{W}}(E_0) = 0$ . Then, we also have that  $\mathbb{P}_{\mathbf{W}}^m(E_0^m) = 0$  and thus

$$\mathbb{P}_{\mathbf{W}}^m((V \cap \mathcal{W}^m) \setminus E_0^m) > 0. \quad (32)$$

This implies that the set  $((\det \circ \eta)^{-1}(\mathbb{R} \setminus \{0\}) \cap \mathcal{W}^m) \setminus E_0^m$  is not empty, *i.e.*, there exists  $(\bar{\mathbf{W}}^{(1)}, \dots, \bar{\mathbf{W}}^{(m)}) \in \mathcal{W}^m \setminus E_0^m$  such that the rows  $\bar{\mathbf{W}}_{i_1, :}^{(1)}, \dots, \bar{\mathbf{W}}_{i_m, :}^{(m)}$  are linearly independent. Since the measure zero set  $E_0$  was arbitrary, this concludes the proof.  $\square$

## B.2. Proof of Theorem 3.1

This section presents the main results building up to Theorem 3.1.

For all  $\mathbf{W} \in \mathcal{W}$ , we are going to denote by  $\hat{\mathbf{W}}^{(\mathbf{W})}$  some estimator of  $\mathbf{W}$ . The following result provides conditions under which if  $\hat{\mathbf{W}}^{(\mathbf{W})}$  allows a perfect fit of the ground-truth distribution  $p(y \mid \mathbf{x}, \mathbf{W})$ , then the representation  $\mathbf{f}_\theta$  and the parameter  $\mathbf{W}$  are identified up to an invertible linear transformation. Many works have showed similar results in various context (Hyvärinen & Morioka, 2016; Khemakhem et al., 2020a; Roeder et al., 2021; Ahuja et al., 2022c). We reuse some of their proof techniques.

**Theorem B.4** (Linear identifiability). *Let  $\hat{\mathbf{W}}^{(\cdot)} : \mathcal{W} \rightarrow \mathbb{R}^{k \times m}$ . Suppose Assumptions 3.2 to 3.4 hold and that, for  $\mathbb{P}_{\mathbf{W}}$ -almost every  $\mathbf{W} \in \mathcal{W}$  and all  $\mathbf{x} \in \mathcal{X}$ , the following holds*

$$\text{KL}(p(y; \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_\theta(\mathbf{x})) \parallel p(y; \mathbf{W} \mathbf{f}_\theta(\mathbf{x}))) = 0. \quad (33)$$

*Then, there exists an invertible matrix  $\mathbf{L} \in \mathbb{R}^{m \times m}$  such that, for all  $\mathbf{x} \in \mathcal{X}$ ,  $\mathbf{f}_\theta(\mathbf{x}) = \mathbf{L} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$  and such that, for  $\mathbb{P}_{\mathbf{W}}$ -almost every  $\mathbf{W} \in \mathcal{W}$ ,  $\hat{\mathbf{W}}^{(\mathbf{W})} = \mathbf{W} \mathbf{L}$*

*Proof.* By Assumption 3.2, Equation (33) implies that, for  $\mathbb{P}_{\mathbf{W}}$ -almost every  $\mathbf{W}$  and all  $\mathbf{x} \in \mathcal{X}$ ,  $\mathbf{W} \mathbf{f}_\theta(\mathbf{x}) = \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$ .

Assumption 3.4 combined with Lemma B.3 ensures that we can construct an invertible matrix  $\mathbf{U} := \begin{bmatrix} \mathbf{W}_{i_1, :}^{(1)} \\ \vdots \\ \mathbf{W}_{i_{d_z}, :}^{(d_z)} \end{bmatrix}$  such

that  $\mathbf{U} \mathbf{f}_\theta(\mathbf{x}) = \hat{\mathbf{U}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$  for all  $\mathbf{x} \in \mathcal{X}$  where  $\hat{\mathbf{U}} := \begin{bmatrix} \hat{\mathbf{W}}_{i_1, :}^{(\mathbf{W}^{(1)})} \\ \vdots \\ \hat{\mathbf{W}}_{i_{d_z}, :}^{(\mathbf{W}^{(d_z)})} \end{bmatrix}$ . Left-multiplying by  $\mathbf{U}^{-1}$  on both sides yields

$\mathbf{f}_\theta(\mathbf{x}) = \mathbf{L} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$ , where  $\mathbf{L} := \mathbf{U}^{-1} \hat{\mathbf{U}}$ . Using the invertible matrix  $\mathbf{F}$  from Assumption 3.3, we can thus write  $\mathbf{F} = \mathbf{L} \hat{\mathbf{F}}$  where we defined  $\hat{\mathbf{F}} := [\mathbf{f}_{\hat{\theta}}(\mathbf{x}^{(1)}), \dots, \mathbf{f}_{\hat{\theta}}(\mathbf{x}^{(d_z)})]$ . Since  $\mathbf{F}$  is invertible, so are  $\mathbf{L}$  and  $\hat{\mathbf{F}}$ .

By substituting  $\mathbf{F} = \mathbf{L} \hat{\mathbf{F}}$  in  $\mathbf{W} \mathbf{F} = \hat{\mathbf{W}}^{(\mathbf{W})} \hat{\mathbf{F}}$ , we obtain  $\mathbf{W} \mathbf{L} \hat{\mathbf{F}} = \hat{\mathbf{W}}^{(\mathbf{W})} \hat{\mathbf{F}}$ . By right-multiplying both sides by  $\hat{\mathbf{F}}^{-1}$ , we obtain  $\mathbf{W} \mathbf{L} = \hat{\mathbf{W}}^{(\mathbf{W})}$ .  $\square$

The following theorem is where most of the theoretical contribution of this work lies. Note that Theorem 3.1, from the main text, is a straightforward application of this result.

**Theorem B.5.** (Disentanglement via task sparsity) *Let  $\hat{\mathbf{W}}^{(\cdot)} : \mathcal{W} \rightarrow \mathbb{R}^{k \times m}$ . Suppose Assumptions 3.2 to 3.6 hold and that, for  $\mathbb{P}_{\mathbf{W}}$ -almost every  $\mathbf{W} \in \mathcal{W}$  and all  $\mathbf{x} \in \mathcal{X}$ , the following holds*

$$\text{KL}(p(y; \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_\theta(\mathbf{x})) \parallel p(y; \mathbf{W} \mathbf{f}_\theta(\mathbf{x}))) = 0. \quad (34)$$

*Moreover, assume that  $\mathbb{E} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \mathbb{E} \|\mathbf{W}\|_{2,0}$ , where both expectations are taken w.r.t.  $\mathbb{P}_{\mathbf{W}}$  and  $\|\mathbf{W}\|_{2,0} := \sum_{j=1}^m \mathbb{1}(\mathbf{W}_{:,j} \neq \mathbf{0})$  with  $\mathbb{1}(\cdot)$  the indicator function. Then,  $\mathbf{f}_\theta$  is disentangled w.r.t.  $\mathbf{f}_\theta$  (Definition 1.1).*

*Proof.* First of all, by Assumptions 3.2 to 3.4, we can apply Theorem B.4 to conclude that  $\mathbf{f}_\theta(\mathbf{x}) = \mathbf{L} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$  and  $\mathbf{W} \mathbf{L} = \hat{\mathbf{W}}^{(\mathbf{W})}$  ( $\mathbb{P}_{\mathbf{W}}$ -almost everywhere) for some invertible matrix  $\mathbf{L}$ .

We can thus write  $\mathbb{E} \|\mathbf{W} \mathbf{L}\|_{2,0} \leq \mathbb{E} \|\mathbf{W}\|_{2,0}$ .

We can write

$$\mathbb{E}\|\mathbf{W}\|_{2,0} = \mathbb{E}_{p(S)} \mathbb{E}\left[\sum_{j=1}^m \mathbb{1}(\mathbf{W}_{:j} \neq \mathbf{0}) \mid S\right] \quad (35)$$

$$= \mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{E}[\mathbb{1}(\mathbf{W}_{:j} \neq \mathbf{0}) \mid S] \quad (36)$$

$$= \mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{P}_{\mathbf{W}|S}[\mathbf{W}_{:j} \neq \mathbf{0}] \quad (37)$$

$$= \mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{1}(j \in S), \quad (38)$$

where the last step follows from the definition of  $S$ .

We now perform similar steps for  $\mathbb{E}\|\mathbf{W}\mathbf{L}\|_{2,0}$ :

$$\mathbb{E}\|\mathbf{W}\mathbf{L}\|_{2,0} = \mathbb{E}_{p(S)} \mathbb{E}\left[\sum_{j=1}^m \mathbb{1}(\mathbf{W}\mathbf{L}_{:j} \neq \mathbf{0}) \mid S\right] \quad (39)$$

$$= \mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{E}[\mathbb{1}(\mathbf{W}\mathbf{L}_{:j} \neq \mathbf{0}) \mid S] \quad (40)$$

$$= \mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{P}_{\mathbf{W}|S}[\mathbf{W}\mathbf{L}_{:j} \neq \mathbf{0}] \quad (41)$$

$$= \mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{P}_{\mathbf{W}|S}[\mathbf{W}_{:S}\mathbf{L}_{S,j} \neq \mathbf{0}]. \quad (42)$$

Notice that

$$\mathbb{P}_{\mathbf{W}|S}[\mathbf{W}_{:S}\mathbf{L}_{S,j} \neq \mathbf{0}] = 1 - \mathbb{P}_{\mathbf{W}|S}[\mathbf{W}_{:S}\mathbf{L}_{S,j} = \mathbf{0}] \quad (43)$$

Let  $N_j$  be the support of  $\mathbf{L}_{:j}$ , i.e.,  $N_j := \{i \in [m] \mid \mathbf{L}_{i,j} \neq 0\}$ . When  $S \cap N_j = \emptyset$ ,  $\mathbf{L}_{S,j} = \mathbf{0}$  and thus  $\mathbb{P}_{\mathbf{W}|S}[\mathbf{W}_{:S}\mathbf{L}_{S,j} = \mathbf{0}] = 1$ . When  $S \cap N_j \neq \emptyset$ ,  $\mathbf{L}_{S,j} \neq \mathbf{0}$ , by [Assumption 3.5](#) we have that  $\mathbb{P}_{\mathbf{W}|S}[\mathbf{W}_{:S}\mathbf{L}_{S,j} = \mathbf{0}] = 0$ . Thus

$$\mathbb{P}_{\mathbf{W}|S}[\mathbf{W}_{:S}\mathbf{L}_{S,j} \neq \mathbf{0}] = 1 - \mathbb{1}(S \cap N_j = \emptyset) \quad (44)$$

$$= \mathbb{1}(S \cap N_j \neq \emptyset), \quad (45)$$

which allows us to write

$$\mathbb{E}\|\mathbf{W}\mathbf{L}\|_{2,0} = \mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{1}(S \cap N_j \neq \emptyset). \quad (46)$$

We thus have that

$$\mathbb{E}\|\mathbf{W}\mathbf{L}\|_{2,0} \leq \mathbb{E}\|\mathbf{W}\|_{2,0} \quad (47)$$

$$\mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{1}(S \cap N_j \neq \emptyset) \leq \mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{1}(j \in S). \quad (48)$$

Since  $\mathbf{L}$  is invertible, by [Lemma B.1](#), there exists a permutation  $\sigma : [m] \rightarrow [m]$  such that, for all  $j \in [m]$ ,  $\mathbf{L}_{j,\sigma(j)} \neq 0$ . In other words, for all  $j \in [m]$ ,  $j \in N_{\sigma(j)}$ . Of course we can permute the terms of the l.h.s. of [Equation \(48\)](#), which yields

$$\mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{1}(S \cap N_{\sigma(j)} \neq \emptyset) \leq \mathbb{E}_{p(S)} \sum_{j=1}^m \mathbb{1}(j \in S) \quad (49)$$

$$\mathbb{E}_{p(S)} \sum_{j=1}^m (\mathbb{1}(S \cap N_{\sigma(j)} \neq \emptyset) - \mathbb{1}(j \in S)) \leq 0. \quad (50)$$

We notice that each term  $\mathbb{1}(S \cap N_{\sigma(j)} \neq \emptyset) - \mathbb{1}(j \in S) \geq 0$  since whenever  $j \in S$ , we also have that  $j \in S \cap N_{\sigma(j)}$  (recall  $j \in N_{\sigma(j)}$ ). Thus, the l.h.s. of Equation (50) is a sum of non-negative terms which is itself non-positive. This means that every term in the sum is zero:

$$\forall S \in \mathcal{S}, \forall j \in [m], \mathbb{1}(S \cap N_{\sigma(j)} \neq \emptyset) = \mathbb{1}(j \in S). \quad (51)$$

Importantly,

$$\forall j \in [m], \forall S \in \mathcal{S}, j \notin S \implies S \cap N_{\sigma(j)} = \emptyset, \quad (52)$$

and since  $S \cap N_{\sigma(j)} = \emptyset \iff N_{\sigma(j)} \subseteq S^c$  we have that

$$\forall j \in [m], \forall S \in \mathcal{S}, j \notin S \implies N_{\sigma(j)} \subseteq S^c \quad (53)$$

$$\forall j \in [m], N_{\sigma(j)} \subseteq \bigcap_{S \in \mathcal{S} | j \notin S} S^c. \quad (54)$$

By Assumption 3.6, we have that  $\bigcup_{S \in \mathcal{S} | j \notin S} S = [m] \setminus \{j\}$ . By taking the complement on both sides and using De Morgan's law, we get  $\bigcap_{S \in \mathcal{S} | j \notin S} S^c = \{j\}$ , which implies that  $N_{\sigma(j)} = \{j\}$  by Equation (54). Thus,  $L = DP$  where  $D$  is an invertible diagonal matrix and  $P$  is a permutation matrix.  $\square$

Before presenting Theorem 3.1 from the main text, we first present a variation of it where we constrain  $\mathbb{E}\|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0}$  to be smaller than  $\mathbb{E}\|\mathbf{W}\|_{2,0}$ . We note that this is weaker than imposing  $\|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \|\mathbf{W}\|_{2,0}$  for all  $\mathbf{W} \in \mathcal{W}$ , as is the case in Problem (4) of Theorem 3.1. Note that Appendix B.3 presents a natural relaxation of Problem (55) which we experiment with in Appendix D.2.5.

**Theorem B.6** (Sparse multitask learning for disentanglement). *Let  $\hat{\theta}$  be a minimizer of*

$$\begin{aligned} & \min_{\hat{\theta}} \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \\ & \text{s.t. } \forall \mathbf{W} \in \mathcal{W}, \hat{\mathbf{W}}^{(\mathbf{W})} \in \arg \min_{\tilde{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \\ & \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\mathbf{W}\|_{2,0}, \end{aligned} \quad (55)$$

where  $\mathbb{P}_{\mathbf{W}}$  and  $p(\mathbf{x}, y | \mathbf{W})$  are described in Section 3.1. Under Assumptions 3.2 to 3.6 and if  $\mathbf{f}_{\hat{\theta}}$  is continuous for all  $\hat{\theta}$ ,  $\mathbf{f}_{\hat{\theta}}$  is disentangled w.r.t.  $\mathbf{f}_{\theta}$  (Definition 1.1).

*Proof.* First, notice that

$$0 \leq \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x} | \mathbf{W})} \text{KL}(p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) || p(y; \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x}))) \quad (56)$$

$$\mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) \leq \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})). \quad (57)$$

This means the objective is minimized (without constraint) if and only if

$$\mathbb{E}_{p(\mathbf{x} | \mathbf{W})} \text{KL}(p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) || p(y; \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x}))) = 0 \quad (58)$$

$\mathbb{P}_{\mathbf{W}}$ -almost everywhere. For a fixed  $\mathbf{W}$ , this equality holds if and only if the KL equals zero  $p(\mathbf{x} | \mathbf{W})$ -almost everywhere, which, by Assumption 3.2, is true if and only if  $\mathbf{W} \mathbf{f}_{\theta}(\mathbf{x}) = \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$   $p(\mathbf{x} | \mathbf{W})$ -almost everywhere. Since both  $\mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})$  and  $\hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$  are continuous functions of  $\mathbf{x}$ , the equality holds over  $\mathcal{X}$  (the support of  $p(\mathbf{x} | \mathbf{W})$ ).

This unconstrained global minimum can actually be achieved by respecting the constraints of Problem (55) simply by setting  $\hat{\theta} := \theta$  and  $\hat{\mathbf{W}}^{(\mathbf{W})} := \mathbf{W}$ . Indeed, the first constraint is satisfied because, for all  $\mathbf{W}$ ,

$$0 \leq \mathbb{E}_{p(\mathbf{x} | \mathbf{W})} \text{KL}(p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) || p(y; \tilde{\mathbf{W}} \mathbf{f}_{\theta}(\mathbf{x}))) \quad (59)$$

$$\mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \mathbf{W} \mathbf{f}_{\theta}(\mathbf{x})) \leq \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\theta}(\mathbf{x})), \quad (60)$$

and clearly the lower bound is attained when  $\tilde{\mathbf{W}} := \mathbf{W}$ . The second constraint is trivially satisfied, since  $\mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} = \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\mathbf{W}\|_{2,0}$ .

The above implies that if  $\hat{\theta}$  is some minimizer of Problem (55), we must have that, (i) for  $\mathbb{P}_{\mathbf{W}}$ -almost every  $\mathbf{W}$ ,  $\mathbf{W} \mathbf{f}_{\theta}(\mathbf{x}) = \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$  for all  $\mathbf{x} \in \mathcal{X}$ , (ii)  $\mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_0 \leq \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\mathbf{W}\|_0$ . Thus, Theorem B.5 implies the desired conclusion.  $\square$

Based on [Theorem B.6](#), we can slightly adjust the argument to prove [Theorem 3.1](#) from the main text.

**Theorem 3.1** (Sparse multi-task learning for disentanglement). *Let  $\hat{\theta}$  be a minimizer of*

$$\begin{aligned} \min_{\hat{\theta}} \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) \\ \text{s.t. } \hat{\mathbf{W}}^{(\mathbf{W})} \in \arg \min_{\tilde{\mathbf{W}} \text{ s.t. } \|\tilde{\mathbf{W}}\|_{2,0} \leq \|\mathbf{W}\|_{2,0}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) , \end{aligned} \quad (4)$$

where the constraint holds for all  $\mathbf{W} \in \mathcal{W}$  and where  $\mathbb{P}_{\mathbf{W}}$  and  $p(\mathbf{x}, y | \mathbf{W})$  are described in [Section 3.1](#). Under [Assumptions 3.2 to 3.6](#) and if  $\mathbf{f}_{\hat{\theta}}$  is continuous for all  $\hat{\theta}$ ,  $\mathbf{f}_{\hat{\theta}}$  is disentangled w.r.t.  $\mathbf{f}_{\theta}$  ([Definition 1.1](#)).

*Proof.* The first part of the argument in the proof of [Theorem B.6](#) applies here as well, meaning: unconstrained minimization of the objective holds if and only if, for  $\mathbb{P}_{\mathbf{W}}$ -almost every  $\mathbf{W}$  and all  $\mathbf{x} \in \mathcal{X}$ ,  $\mathbf{W} \mathbf{f}_{\theta}(\mathbf{x}) = \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$ . And again, this unconstrained minimum can be achieved by respecting the constraint of [Problem \(4\)](#) simply by setting  $\hat{\theta} := \theta$  and  $\hat{\mathbf{W}}^{(\mathbf{W})} := \mathbf{W}$ .

This means that if  $\hat{\theta}$  is some minimizer of [Problem \(4\)](#), we must have (i) for  $\mathbb{P}_{\mathbf{W}}$ -almost every  $\mathbf{W}$ ,  $\mathbf{W} \mathbf{f}_{\theta}(\mathbf{x}) = \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})$  for all  $\mathbf{x} \in \mathcal{X}$  and (ii) for all  $\mathbf{W} \in \mathcal{W}$ ,  $\|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \|\mathbf{W}\|_{2,0}$ . Of course the latter point implies  $\mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\mathbf{W}\|_{2,0}$ , which allows us to apply [Theorem B.5](#) to obtain the desired conclusion.  $\square$

### B.3. Regularization in the outer problem instead of in the inner problem

[Theorem B.6](#) presented an alternative bilevel optimization problem to the one of [Theorem 3.1](#) in the main text. Essentially, the difference is that the constraints  $\|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \|\mathbf{W}\|_{2,0}$  for all  $\mathbf{W} \in \mathcal{W}$  are replaced by the unique constraint  $\mathbb{E} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \mathbb{E} \|\mathbf{W}\|_{2,0}$ , which is a weaker constraint.

In [Section 3.4](#), we introduced a tractable relaxation of the problem of [Theorem 3.1](#). In this section, we introduce a relaxation of the problem of [Theorem B.6](#).

A natural idea is to replace the constraint  $\mathbb{E} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \mathbb{E} \|\mathbf{W}\|_{2,0}$  of [Theorem B.6](#) by a penalty  $\lambda \mathbb{E} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,1}$  in the outer problem, like so:

$$\begin{aligned} \min_{\hat{\theta}} \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \hat{\mathbf{W}}^{(\mathbf{W})} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) + \lambda \mathbb{E}_{\mathbb{P}_{\mathbf{W}}} \|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,1} \\ \text{s.t. } \forall \mathbf{W} \in \mathcal{W}, \hat{\mathbf{W}}^{(\mathbf{W})} \in \arg \min_{\tilde{\mathbf{W}}} \mathbb{E}_{p(\mathbf{x}, y | \mathbf{W})} - \log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})), \end{aligned} \quad (61)$$

in which we can replace the expectations by empirical averages to get

$$\begin{aligned} \min_{\hat{\theta}} \frac{1}{T} \sum_{t=1}^T \left[ -\frac{1}{n} \sum_{(\mathbf{x}, y) \in \mathcal{D}_t} \log p(y; \hat{\mathbf{W}}^{(t)} \mathbf{f}_{\hat{\theta}}(\mathbf{x})) + \lambda \|\hat{\mathbf{W}}^{(t)}\|_{2,1} \right] \\ \text{s.t. } \hat{\mathbf{W}}^{(t)} \in \arg \min_{\tilde{\mathbf{W}}} \frac{1}{n} \sum_{(\mathbf{x}, y) \in \mathcal{D}_t} -\log p(y; \tilde{\mathbf{W}} \mathbf{f}_{\hat{\theta}}(\mathbf{x})). \end{aligned} \quad (62)$$

This can be optimized in the same way as [Problem \(6\)](#) via implicit differentiation and standard gradient descent algorithms. The essential difference between [Problem \(62\)](#) and [Problem \(6\)](#) is that the former has regularization in the outer problem instead of in the inner problem. From a practical point of view, this problem is typically simpler than [Problem \(6\)](#) since the inner objective is generally smooth, and standard implicit differentiation techniques apply (the non-smooth term  $\|\tilde{\mathbf{W}}\|_{2,1}$  in the inner objective of [Problem \(6\)](#) requiring some care with implicit differentiation; [Bertrand et al., 2022](#)). We provide some experimental results in [Appendix D.2.5](#) demonstrating that this alternative works as well.

### B.4. What can go wrong when [Assumption 3.5](#) is violated?

[Theorem B.4](#) allowed us to conclude that  $\hat{\mathbf{W}}^{(\mathbf{W})} = \mathbf{W} \mathbf{L}$  for  $\mathbb{P}_{\mathbf{W}}$ -almost every  $\mathbf{W}$  and that  $\mathbf{L} \mathbf{f}_{\hat{\theta}}(\mathbf{x}) = \mathbf{f}_{\theta}(\mathbf{x})$  for all  $\mathbf{x} \in \mathcal{X}$ . The rest of the argument leading up to [Theorem 3.1](#) essentially amounts to showing that having  $\|\hat{\mathbf{W}}^{(\mathbf{W})}\|_{2,0} \leq \|\mathbf{W}\|_{2,0}$

for all  $\mathbf{W} \in \mathcal{W}$  forces  $\mathbf{L}$  to be a permutation-scaling matrix. The intuition is that  $\|\mathbf{W}\mathbf{L}\|_{2,0} \leq \|\mathbf{W}\|_{2,0}$  everywhere should force  $\mathbf{L}$  to be sparse, and maximal sparsity is precisely when  $\mathbf{L}$  is a permutation-scaling matrix. But just how many  $\mathbf{W}$  do we need and how diverse should they be to make this argument formal? Our answer is given by [Assumption 3.5](#). But what can go wrong when this assumption is not satisfied? To answer this question, we construct a counterexample in which the distribution  $\mathbb{P}_{\mathbf{W}}$  satisfies [Assumption 3.6](#) but not [Assumption 3.5](#) and a matrix  $\mathbf{L}$  that satisfies the constraint  $\|\mathbf{W}\mathbf{L}\|_{2,0} \leq \|\mathbf{W}\|_{2,0}$  everywhere but that is not a permutation-scaling matrix. Consider a distribution  $\mathbb{P}_{\mathbf{W}}$  with support  $\mathcal{W} := \{[1, 1, 0], [1, 0, 1], [0, 1, 1]\}$  (which is finite) and let

$$\mathbf{L} := \begin{bmatrix} 3 & -1 & -1 \\ -1 & 1 & 3 \\ 1 & 3 & 1 \end{bmatrix}, \quad (63)$$

which, of course, is not a permutation-scaling matrix. One can then compute to show that the sparsity constraint holds for all  $\mathbf{W} \in \mathcal{W}$ :

$$\|[1\ 1\ 0]\mathbf{L}\|_{2,0} = \|[2\ 0\ 2]\|_{2,0} \leq 2 = \|[1\ 1\ 0]\|_{2,0} \quad (64)$$

$$\|[1\ 0\ 1]\mathbf{L}\|_{2,0} = \|[4\ 2\ 0]\|_{2,0} \leq 2 = \|[1\ 0\ 1]\|_{2,0} \quad (65)$$

$$\|[0\ 1\ 1]\mathbf{L}\|_{2,0} = \|[0\ 4\ 4]\|_{2,0} \leq 2 = \|[0\ 1\ 1]\|_{2,0}. \quad (66)$$

This means that, with such a  $\mathbb{P}_{\mathbf{W}}$ , solving the bilevel problem of [Theorem 3.1](#) will not necessarily lead to a disentangled representation since one could fall on a “bad”  $\mathbf{L}$  such as the one defined above.

### B.5. Assumption 3.6 holds with high probability when the number of supports is large

In this section, we provide a probabilistic argument showing that [Assumption 3.6](#) holds with high probability when the number of supports is large. Let  $\mathcal{S}^{(T)} := \{S^{(1)}, S^{(2)}, \dots, S^{(T)}\}$  be the set of supports observed, where  $T$  is the number of supports. To make this argument, we will assume that the  $S^{(t)}$  are sampled independently and identically. Moreover,  $\mathbb{P}[i \in S^{(t)}] = p \in (0, 1)$  and these events are assumed independent.

The next proposition shows that the probability that [Assumption 3.6](#) fails under the above model is very small when  $T$  is large.

**Proposition B.7.** *Given the probabilistic model described above, we have*

$$\mathbb{P} \left[ \exists j \in [m] \text{ s.t. } \bigcup_{S \in \mathcal{S}^{(T)} | j \notin S} S \neq [m] \setminus \{j\} \right] \leq m(m-1)(1-p(1-p))^T \xrightarrow{T \rightarrow \infty} 0. \quad (67)$$

*Proof.* By rewriting slightly the original probability statement and applying the union bound, we get

$$\mathbb{P} \left[ \exists j \in [m] \text{ s.t. } \bigcup_{S \in \mathcal{S}^{(T)} | j \notin S} S \neq [m] \setminus \{j\} \right] \quad (68)$$

$$= \mathbb{P} \left[ \exists j \in [m], i \in [m] \setminus \{j\} \text{ s.t. } i \notin \bigcup_{S \in \mathcal{S}^{(T)} | j \notin S} S \right] \quad (69)$$

$$\leq \sum_{j=1}^m \sum_{i \in [m] \setminus \{j\}} \mathbb{P} \left[ i \notin \bigcup_{S \in \mathcal{S}^{(T)} | j \notin S} S \right], \quad (70)$$

We can further write

$$\mathbb{P} \left[ i \notin \bigcup_{S \in \mathcal{S}^{(T)} | j \notin S} S \right] = \mathbb{P} \left[ \forall t \in [T], j \notin S^{(t)} \implies i \notin S^{(t)} \right] \quad (71)$$

$$= \mathbb{P} \left[ \forall t \in [T], j \in S^{(t)} \vee i \notin S^{(t)} \right] \quad (72)$$

$$= \prod_{t=1}^T \mathbb{P} \left[ j \in S^{(t)} \vee i \notin S^{(t)} \right], \quad (73)$$

where the last step holds because the supports  $S^{(t)}$  are mutually independent. We continue and get

$$\mathbb{P} \left[ i \notin \bigcup_{S \in \mathcal{S}^{(T)} | j \notin S} S \right] = \prod_{t=1}^T \mathbb{P} \left[ j \in S^{(t)} \vee i \notin S^{(t)} \right] \quad (74)$$

$$= \prod_{t=1}^T (1 - \mathbb{P} [j \notin S^{(t)} \wedge i \in S^{(t)}]) \quad (75)$$

$$= \prod_{t=1}^T (1 - \mathbb{P} [j \notin S^{(t)}] \mathbb{P} [i \in S^{(t)}]) \quad (76)$$

$$= \prod_{t=1}^T (1 - (1-p)p), \quad (77)$$

where we used the fact that the events  $j \notin S^{(t)}$  and  $i \in S^{(t)}$  are independent (when  $i \neq j$ ). Bringing everything together, one gets

$$\mathbb{P} \left[ \exists j \in [m] \text{ s.t. } \bigcup_{S \in \mathcal{S}^{(T)} | j \notin S} S \neq [m] \setminus \{j\} \right] \leq \sum_{j=1}^m \sum_{i \in [m] \setminus \{j\}} \prod_{t=1}^T (1 - (1-p)p) \quad (78)$$

$$= m(m-1)(1 - (1-p)p)^T \quad (79)$$

$$(80)$$

which converges to 0 when  $T \rightarrow \infty$  since  $0 < 1 - (1-p)p < 1$ .  $\square$

## B.6. A distribution without density satisfying Assumption 3.5

Interestingly, there are distributions over  $\mathbf{W}_{1,S} \mid S$  that do not have a density w.r.t. the Lebesgue measure, but still satisfy Assumption 3.5. This is the case, e.g., when  $\mathbf{W}_{1,S} \mid S$  puts uniform mass over a  $(|S| - 1)$ -dimensional sphere embedded in  $\mathbb{R}^{|S|}$  and centered at zero. In that case, for all  $\mathbf{a} \in \mathbb{R}^{|S|} \setminus \{0\}$ , the intersection of  $\text{span}\{\mathbf{a}\}^\perp$ , which is  $(|S| - 1)$ -dimensional, with the  $(|S| - 1)$ -dimensional sphere is  $(|S| - 2)$ -dimensional and thus has probability zero of occurring. One can certainly construct more exotic examples of measures satisfying Assumption 3.5 that concentrate mass on lower dimensional manifolds.

## C. Optimization details

### C.1. Group Lasso SVM Dual

**Notation.** The Fenchel conjugate of a function  $h : \mathbb{R}^d \rightarrow \mathbb{R}$  is written  $h^*$  and is defined for any  $y \in \mathbb{R}^d$ , by  $h^*(y) = \sup_{x \in \mathbb{R}^d} \langle x, y \rangle - h(x)$ .

**Definition C.1.** (*Primal Group Lasso Soft-Margin Multiclass SVM.*) The primal problem of the group Lasso soft-margin multiclass SVM is defined as

$$\min_{\mathbf{W} \in \mathbb{R}^{k \times m}} \mathcal{L}_{\text{in}}(\mathbf{W}; \mathbf{F}, \mathbf{Y}) := \sum_{i=1}^n \max_{l \in [k]} (1 + (\mathbf{W}_{y_i} - \mathbf{W}_l) \mathbf{F}_i - \mathbf{Y}_{il}) + \lambda_1 \|\mathbf{W}\|_{2,1} + \frac{\lambda_2}{2} \|\mathbf{W}\|^2 \quad (81)$$

**Proposition C.2.** (Dual Group Lasso Soft-Margin Multiclass SVM.) *The dual of the inner problem with  $\mathcal{L}_{\text{in}}$  as defined in (8) writes*

$$\begin{aligned} \min_{\mathbf{\Lambda} \in \mathbb{R}^{n \times k}} \quad & \frac{1}{\lambda_2} \sum_{j=1}^m \|\text{BST}((\mathbf{Y} - \mathbf{\Lambda})^\top \mathbf{F}_{:j}, \lambda_1)\|^2 + \langle \mathbf{Y}, \mathbf{\Lambda} \rangle \\ \text{s.t.} \quad & \forall i, l, \in [n] \times [k], \sum_{l'=1}^k \mathbf{\Lambda}_{il'} = 1 \text{ and } \mathbf{\Lambda}_{il} \geq 0, \end{aligned} \quad (10)$$

with  $\text{BST} : (\mathbf{a}, \tau) \mapsto (1 - \tau/\|\mathbf{a}\|)_+ \mathbf{a}$  is the block soft-thresholding operator;  $\mathbf{F} \in \mathbb{R}^{n \times m}$  the concatenation of  $\{\mathbf{f}_{\hat{\theta}}(x)\}_{(x,y) \in \mathcal{D}^{\text{train}}}$ . In addition, the primal-dual link writes,  $\forall j \in [m]$ ,  $\mathbf{W}_{:j} = \text{BST}((\mathbf{Y} - \mathbf{\Lambda})^\top \mathbf{F}_{:j}, \lambda_1) / \lambda_2$ .

The primal objective 81 can be hard to minimize with modern solvers. Moreover in few-shot learning applications, the number of features  $m$  is usually much larger than the number of samples  $n$  (in Lee et al. 2019,  $m = 1.6 \cdot 10^4$  and  $n \leq 25$ ), hence we solve the dual of Problem (81).

*Proof of Proposition C.2.* Let  $g : \mathbf{u} \mapsto \lambda_1 \|\mathbf{u}\| + \frac{\lambda_2}{2} \|\mathbf{u}\|^2$ . Proof of Proposition C.2 is composed of the following lemmas.

**Lemma C.3.** i) *The dual of Problem (81) is*

$$\begin{aligned} \min_{\mathbf{\Lambda} \in \mathbb{R}^{n \times k}} \quad & \sum_{j=1}^m g^*((\mathbf{Y} - \mathbf{\Lambda})^\top \mathbf{F}_{:j}) + \langle \mathbf{Y}, \mathbf{\Lambda} \rangle \\ \text{s.t.} \quad & \forall i \in [n], \sum_{l=1}^k \mathbf{\Lambda}_{il} = 1, \quad \forall i \in [n], l \in [k], \mathbf{\Lambda}_{il} \geq 0, \end{aligned} \quad (82)$$

where  $g^*$  is the Fenchel conjugate of the function  $g$ .

ii) *The Fenchel conjugate of the function  $g$  writes*

$$\forall \mathbf{v} \in \mathbb{R}^K, g^*(\mathbf{v}) = \frac{1}{\lambda_2} \|\text{BST}(\mathbf{v}, \lambda_1)\|^2. \quad (83)$$

Lemmas C.3 i) and C.3 ii) yields Proposition C.2.

*Proof of Lemma C.3 i).* The Lagrangian of Problem (81) writes:

$$\mathcal{L}(\mathbf{W}, \boldsymbol{\xi}, \mathbf{\Lambda}) = \sum_{j=1}^m g(\mathbf{W}_{:j}) + \sum_i \xi_i + \sum_{i=1}^n \sum_{l=1}^k (1 - \xi_i - \mathbf{W}_{\mathbf{y}_i} \cdot \mathbf{F}_i + \mathbf{W}_l \cdot \mathbf{F}_i - \mathbf{Y}_{il}) \mathbf{\Lambda}_{il}. \quad (84)$$

$\partial_{\boldsymbol{\xi}} \mathcal{L}(\mathbf{W}, \boldsymbol{\xi}, \mathbf{\Lambda}) = 0$  yields  $\forall i \in [n], \sum_{l=1}^k \mathbf{\Lambda}_{il} = 1$ . Then the Lagrangian rewrites

$$\begin{aligned} \min_{\mathbf{W}} \min_{\boldsymbol{\xi}} \mathcal{L}(\mathbf{W}, \boldsymbol{\xi}, \mathbf{\Lambda}) &= \min_{\mathbf{W}, \boldsymbol{\xi}} \sum_{j=1}^m g(\mathbf{W}_{:j}) + \sum_{i=1}^n \xi_i + \sum_{i=1}^n \sum_{l=1}^k (-\xi_i - \mathbf{W}_{\mathbf{y}_i} \cdot \mathbf{F}_i + \mathbf{W}_l \cdot \mathbf{F}_i - \mathbf{Y}_{il}) \mathbf{\Lambda}_{il} \\ &= \sum_{j=1}^m \min_{\mathbf{W}_{:j}} g(\mathbf{W}_{:j}) - \underbrace{\sum_{i=1}^n \sum_{l=1}^k (\mathbf{F}_i \cdot \mathbf{Y}_{il} - \mathbf{F}_i \cdot \mathbf{\Lambda}_{il}) \mathbf{W}_l}_{= \langle (\mathbf{Y} - \mathbf{\Lambda})^\top \mathbf{F}_{:j}, \mathbf{W}_{:j} \rangle} - \sum_{i=1}^n \sum_{l=1}^k \mathbf{Y}_{il} \mathbf{\Lambda}_{il} \\ &= \underbrace{\sum_{j=1}^m \min_{\mathbf{W}_{:j}} g(\mathbf{W}_{:j}) - \sum_{i=1}^n \sum_{l=1}^k (\mathbf{F}_i \cdot \mathbf{Y}_{il} - \mathbf{F}_i \cdot \mathbf{\Lambda}_{il}) \mathbf{W}_l}_{= -g^*((\mathbf{Y} - \mathbf{\Lambda})^\top \mathbf{F}_{:j})} - \sum_{i=1}^n \sum_{l=1}^k \mathbf{Y}_{il} \mathbf{\Lambda}_{il}. \end{aligned}$$

Then the dual problem writes:

$$\min_{\mathbf{\Lambda} \in \mathbb{R}^{n \times k}} \sum_{j=1}^m g^* \left( (\mathbf{Y} - \mathbf{\Lambda})^\top \mathbf{F}_{:j} \right) + \langle \mathbf{Y}, \mathbf{\Lambda} \rangle \quad (85)$$

$$\text{s. t. } \forall i \in [n] \quad \sum_{l=1}^k \mathbf{\Lambda}_{il} = 1, \forall i \in [n], l \in [k], \quad \mathbf{\Lambda}_{il} \geq 0. \quad (86)$$

□

*Proof of Lemma C.3 ii).* Let  $h : \mathbf{u} \mapsto \|\mathbf{u}\|_2 + \frac{\kappa}{2} \|\mathbf{u}\|_2^2$ . The proof of Lemma C.3 i) is done using the following steps.

**Lemma C.4.** i)  $h^*(\mathbf{v}) = \frac{1}{2\kappa} \|\mathbf{v}\|_2^2 - \left( \frac{\kappa}{2} \|\cdot\|_2^2 \square \|\cdot\|_2 \right) (\mathbf{v}/\kappa)$ .

ii)  $\left( \frac{\kappa}{2} \|\cdot\|_2^2 \square \|\cdot\|_2 \right) (\mathbf{v}) = \frac{\kappa}{2} \|\mathbf{v}\|_2^2 - \frac{1}{2\kappa} \|\text{BST}(\kappa\mathbf{v}, 1)\|_2^2$ .

*Proof of Lemma C.4 i).* With  $\kappa = \lambda_2/\lambda_1$ , the Fenchel transform of  $h : \mathbf{w} \mapsto \|\mathbf{w}\|_2 + \kappa \|\mathbf{w}\|_2^2$ .

$$\begin{aligned} h(\mathbf{u}) &= \|\mathbf{u}\|_2 + \frac{\kappa}{2} \|\mathbf{u}\|_2^2 \\ h^*(\mathbf{v}) &= \sup_{\mathbf{w}} \left( \mathbf{v}^\top \mathbf{w} - \|\mathbf{w}\|_2 - \frac{\kappa}{2} \|\mathbf{w}\|_2^2 \right) \\ &= \frac{1}{2\kappa} \|\mathbf{v}\|_2^2 + \sup_{\mathbf{w}} \left( -\frac{\kappa}{2} \|\mathbf{w} - \mathbf{v}/\kappa\|_2^2 - \|\mathbf{w}\|_2 \right) \\ &= \frac{1}{2\kappa} \|\mathbf{v}\|_2^2 - \inf_{\mathbf{w}} \left( \frac{\kappa}{2} \|\mathbf{w} - \mathbf{v}/\kappa\|_2^2 + \|\mathbf{w}\|_2 \right) \\ &= \frac{1}{2\kappa} \|\mathbf{v}\|_2^2 - \left( \frac{\kappa}{2} \|\cdot\|_2^2 \square \|\cdot\|_2 \right) (\mathbf{v}/\kappa). \end{aligned}$$

□

*Proof of Lemma C.4 ii).*

$$\begin{aligned} \left( \frac{\kappa}{2} \|\cdot\|_2^2 \square \|\cdot\|_2 \right) (\mathbf{v}) &= \left( \frac{\kappa}{2} \|\cdot\|_2^2 \square \|\cdot\|_2 \right)^{**} (\mathbf{v}) \\ &= \left( \frac{1}{2\kappa} \|\cdot\|_2^2 + \iota_{\mathcal{B}_2} \right)^* (\mathbf{v}) \\ &= \sup_{\|\mathbf{w}\|_2 \leq 1} \left( \mathbf{v}^\top \mathbf{w} - \frac{1}{2\kappa} \|\mathbf{w}\|_2^2 \right) \\ &= \frac{\kappa}{2} \|\mathbf{v}\|_2^2 + \sup_{\|\mathbf{w}\|_2 \leq 1} -\frac{1}{2\kappa} \|\kappa\mathbf{v} - \mathbf{w}\|_2^2 \\ &= \frac{\kappa}{2} \|\mathbf{v}\|_2^2 - \frac{1}{2\kappa} \|\text{BST}(\kappa\mathbf{v}, 1)\|_2^2. \end{aligned}$$

□

$$\begin{aligned} g^*(\mathbf{u}) &= \lambda_1 h^*(\mathbf{u}/\lambda_1) \\ &= \frac{\lambda_1}{2\kappa} \|\text{BST}(\mathbf{u}/\lambda_1, 1)\|_2^2 \\ &= \frac{\lambda_1^2}{2\lambda_2} \|\text{BST}(\mathbf{u}/\lambda_1, 1)\|_2^2 \\ &= \frac{1}{\lambda_2} \|\text{BST}(\mathbf{u}, \lambda_1)\|_2^2. \end{aligned}$$

□

□



## D. Experimental details

### D.1. Disentangled representation coupled with sparsity regularization improves generalization

We consider the following data generating process: We sample the ground-truth features  $\mathbf{f}_\theta(\mathbf{x})$  from a Gaussian distribution  $\mathcal{N}(\mathbf{0}, \Sigma)$  where  $\Sigma \in \mathbb{R}^{m \times m}$  and  $\Sigma_{i,j} = 0.9^{|i-j|}$ . Moreover, the labels are given by  $y = \mathbf{w} \cdot \mathbf{f}_\theta(\mathbf{x}) + \epsilon$  where  $\mathbf{w} \in \mathbb{R}^m$ ,  $\epsilon \sim \mathcal{N}(0, 0.04)$  and  $m = 100$ . The ground-truth weight vector  $\mathbf{w}$  is sampled once from  $\mathcal{N}(0, I_{m \times m})$  and mask some of its components to zero: we vary the fraction of meaningful features ( $\ell/m$ ) from very sparse ( $\ell/m = 5\%$ ) to less sparse ( $\ell/m = 80\%$ ) settings. For each case, we study the sample complexity by varying the number of training samples from 25 to 150, but evaluating the generalization performance on a larger test dataset (1000 samples). To generate the entangled representations, we multiply the true latent variables  $\mathbf{f}_\theta(\mathbf{x})$  by a randomly sampled orthogonal matrix  $\mathbf{L}$ , i.e.,  $\mathbf{f}_\theta(\mathbf{x}) := \mathbf{L}\mathbf{f}_\theta(\mathbf{x})$ . For the disentangled representation, we simply consider the true latents, i.e.,  $\mathbf{f}_\theta(\mathbf{x}) := \mathbf{f}_\theta(\mathbf{x})$ . Note that in principle we could have considered an invertible matrix  $\mathbf{L}$  that is not orthogonal for the linearly entangled representation and a component-wise rescaling for the disentangled representation. The advantage of not doing so and opting for our approach is that the conditioning number of the covariance matrix of  $\mathbf{f}_\theta(\mathbf{x})$  is the same for both the entangled and the disentangled, hence offering a fairer comparison.

For both the case of entangled and disentangled representation, we solve the regression problem with Lasso and Ridge regression, where the associated hyperparameters (regularization strength) were inferred using 5-fold cross-validation on the input training dataset. Using both lasso and ridge regression would help us to show the effect of encouraging sparsity.

In Figure 1 for the sparsest case ( $\ell/m = 5\%$ ), we observe that that Disentangled-Lasso approach has the best performance when we have fewer training samples, while the Entangled-Lasso approach performs the worst. As we increase the number of training samples, the performance of Entangled-Lasso approaches that of Disentangled-Lasso, however, learning under the Disentangled-Lasso approach is sample efficient. Disentangled-Lasso obtains  $R^2$  greater than 0.5 with only 25 training samples, while other approaches obtain  $R^2$  close to zero. Also, Disentangled-Lasso converges to the optimal  $R^2$  using only 50 training samples, while Entangled-Lasso does the same with 150 samples.

Note that the improvement due to disentanglement does not happen for the case of ridge regression as expected and there is no difference between the methods Disentangled-Ridge and Entangled-Ridge because the L2 norm is invariant to orthogonal transformation. Also, having sparsity in the underlying task is important. Disentangled-Lasso shows the max improvement for the case of  $\ell/m = 5\%$ , with the gains reducing as we decrease the sparsity in the underlying task ( $\ell/m = 80\%$ ).

### D.2. Disentanglement in 3D Shapes

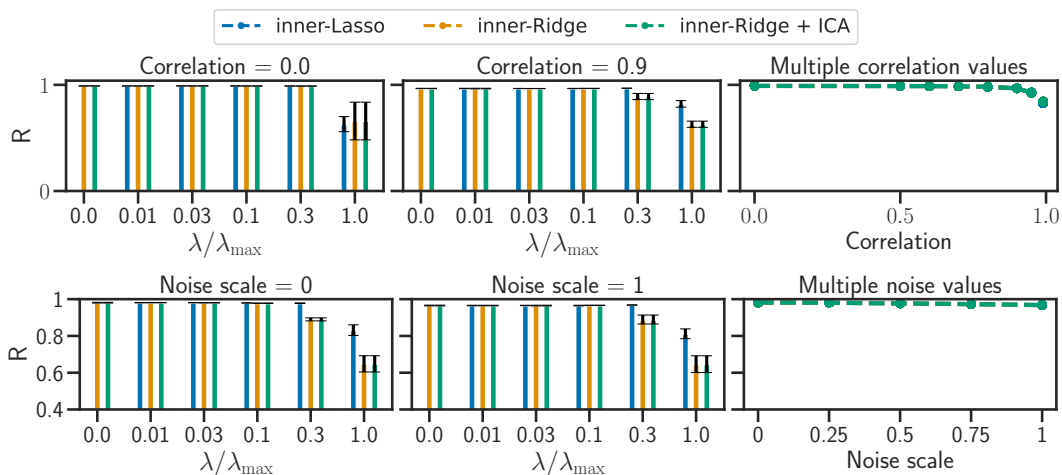


Figure 6. Prediction performance (R Score) for inner-Lasso, inner-Ridge and inner-Ridge combined with ICA as a function of the regularization parameter (left and middle). Varying level of correlation between latents (top) and noise on the latents (bottom). The right columns shows performance of the best hyperparameter for different values of correlation and noise levels.

## D.2.1. DATASET GENERATION

**Details on 3D Shapes.** The 3D Shapes dataset (Burgess & Kim, 2018) contains synthetic images of colored shapes resting in a simple 3D scene. These images vary across 6 factors: Floor hue (10 values linearly spaced in  $[0, 1]$ ); Wall hue (10 values linearly spaced in  $[0, 1]$ ); Object hue (10 values linearly spaced in  $[0, 1]$ ); Scale (8 values linearly spaced in  $[0, 1]$ ); Shape (4 values in  $[0, 1, 2, 3]$ ); and Orientation (15 values linearly spaced in  $[-30, 30]$ ). These are the factors we aim to disentangle. We standardize them to have mean 0 and variance 1. We denote by  $\mathcal{Z} \subset \mathbb{R}^6$ , the set of all possible latent factor combinations. In our framework, this corresponds to the support of the ground-truth features  $\mathbf{f}_\theta(\mathbf{x})$ . We note that the points in  $\mathcal{Z}$  are arranged in a grid-like fashion in  $\mathbb{R}^6$ .

**Task generation.** For all tasks  $t$ , the labelled dataset  $\mathcal{D}_t = \{(\mathbf{x}^{(t,i)}, y^{(t,i)})\}_{i=1}^n$  is generated by first sampling the ground-truth latent variables  $\mathbf{z}^{(t,i)} := \mathbf{f}_\theta(\mathbf{x}^{(t,i)})$  i.i.d. according to some distribution  $p(\mathbf{z})$  over  $\mathcal{Z}$ , while the corresponding input is obtained doing  $\mathbf{x}^{(t,i)} := \mathbf{f}_\theta^{-1}(\mathbf{z}^{(t,i)})$  ( $\mathbf{f}_\theta$  is invertible in 3D Shapes). Then, a sparse weight vector  $\mathbf{w}^{(t)}$  is sampled randomly by doing  $\mathbf{w}^{(t)} := \bar{\mathbf{w}}^{(t)} \odot \mathbf{s}^{(t)}$ , where  $\odot$  is the Hadamard (component-wise) product,  $\bar{\mathbf{w}}^{(t)} \sim \mathcal{N}(\mathbf{0}, I)$  and  $\mathbf{s} \in \{0, 1\}^6$  is a binary vector with independent components sampled from a Bernoulli distribution with ( $p = 0.5$ ). Then, the labels are computed for each example as  $y^{(t,i)} := \mathbf{w}^{(t)} \cdot \mathbf{x}^{(t,i)} + \epsilon^{(t,i)}$ , where  $\epsilon^{(t,i)}$  is independent Gaussian noise. In every task, the dataset has size  $n = 50$ . New tasks are generated continuously as we train. Fig. 4 and 6 explores various choices of  $p(\mathbf{z})$ , i.e., by varying the level of correlation between the latent variables and by varying the level of noise on the ground-truth latents. Fig. 7 shows a visualization of some of these distributions over latents.

**Noise on latents.** To make the dataset slightly more realistic, we get rid of the artificial grid-like structure of the latents by adding noise to it. This procedure transforms  $\mathcal{Z}$  into a new support  $\mathcal{Z}_\alpha$ , where  $\alpha$  is the noise level. Formally,  $\mathcal{Z}_\alpha := \bigcup_{\mathbf{z} \in \mathcal{Z}} \{\mathbf{z} + \mathbf{u}_z\}$  where the  $\mathbf{u}_z$  are i.i.d samples from the uniform over the hypercube

$$\left[-\alpha \frac{\Delta z_1}{2}, \alpha \frac{\Delta z_1}{2}\right] \times \left[-\alpha \frac{\Delta z_2}{2}, \alpha \frac{\Delta z_2}{2}\right] \times \dots \times \left[-\alpha \frac{\Delta z_6}{2}, \alpha \frac{\Delta z_6}{2}\right],$$

where  $\Delta z_i$  denotes the gap between contiguous values of the factor  $z_i$ . When  $\alpha = 0$ , no noise is added and the support  $\mathcal{Z}$  is unchanged, i.e.,  $\mathcal{Z}_1 = \mathcal{Z}$ . As long as  $\alpha \in [0, 1]$ , contiguous points in  $\mathcal{Z}$  cannot be interchanged in  $\mathcal{Z}_\alpha$ . We also clarify that the ground-truth mapping  $\mathbf{f}_\theta$  is modified to  $\mathbf{f}_{\theta,\alpha}$  consequently: for all  $\mathbf{x} \in \mathcal{X}$ ,  $\mathbf{f}_{\theta,\alpha}(\mathbf{x}) := \mathbf{f}_\theta(\mathbf{x}) + \mathbf{u}_z$ . We emphasize that the  $\mathbf{u}_z$  are sampled only once such that  $\mathbf{f}_{\theta,\alpha}(\mathbf{x})$  is actually a deterministic mapping.

**Varying correlations.** To verify that our approach is robust to correlations in the latents, we construct  $p(\mathbf{z})$  as follows: We consider a Gaussian density centred at  $\mathbf{0}$  with covariance  $\Sigma_{i,j} := \rho + \mathbb{1}(i = j)(1 - \rho)$ . Then, we evaluate this density on the points of  $\mathcal{Z}_\alpha$  and renormalize to have a well-defined probability distribution over  $\mathcal{Z}_\alpha$ . We denote by  $p_{\alpha,\rho}(\mathbf{z})$  the distribution obtain by this construction.

In the top rows of Fig. 4 and 6, the latents are sampled from  $p_{\alpha=1,\rho}(\mathbf{z})$  and  $\rho$  varies between 0 and 0.99. In the bottom rows of Fig. 4 and 6, the latents are sampled from  $p_{\alpha,\rho=0.9}(\mathbf{z})$  and  $\alpha$  varies from 0 to 1.

## D.2.2. METRICS

We evaluate disentanglement via the *mean correlation coefficient* (Hyvärinen & Morioka, 2016; Khemakhem et al., 2020a) which is computed as follows: The Pearson correlation matrix  $C$  between the ground-truth features and learned ones is computed. Then,  $\text{MCC} = \max_{\pi \in \text{permutations}} \frac{1}{m} \sum_{j=1}^m |C_{j,\pi(j)}|$ . We also evaluate linear equivalence by performing linear regression to predict the ground-truth factors from the learned ones, and report the mean of the Pearson correlations between the ground-truth latents and the learned ones. This metric is known as the *coefficient of multiple correlations*,  $R$ , and turns out to be the square-root of the more widely known *coefficient of determination*,  $R^2$ . The advantage of using  $R$  over  $R^2$  is that we always have  $\text{MCC} \leq R$ .

## D.2.3. ARCHITECTURE, INNER SOLVER &amp; HYPERPARAMETERS

We use the four-layer convolutional neural network typically used in the disentanglement literature (Locatello et al., 2019). As mentioned in Section 3.4, the norm of the representation  $\mathbf{f}_{\hat{\theta}}(\mathbf{x})$  must be controlled to make sure the regularization remains effective. To do so, we apply batch normalization (Ioffe & Szegedy, 2015) at the very last layer of the neural network and do not learn its scale and shift parameters. Empirically, we do see the expected behavior that, without any normalization, the norm of  $\mathbf{f}_{\hat{\theta}}(\mathbf{x})$  explodes as we train, leading to instabilities and low sparsity.

In these experiments, the distribution  $p(y; \eta)$  used for learning is a Gaussian with fixed variance. In that case, the inner




Figure 7. Visualization of the various distributions over latents. For 4 combinations of correlation levels and noise levels, we show the 2-dimensional histograms of samples from the corresponding distribution over latents described in Appendix D.2.1. Each histogram shows the joint distribution over two latent factors.

problem of Section 3.4 reduces to Lasso regression. Computing the hypergradient w.r.t.  $\theta$  requires solving this inner problem. To do so, we use Proximal Coordinate Descent (Tseng, 2001; Richtárik & Takáč, 2014).

**Details on  $\lambda/\lambda_{\max}$ .** In Fig. 4 and 6, we explore various levels of regularization  $\lambda$ . In our implementation, we set  $\lambda = \epsilon\lambda_{\max}$  where  $\epsilon \geq 0$ . In inner-Lasso, we set  $\lambda_{\max} := \frac{1}{n}\|\mathbf{F}^\top \mathbf{y}\|_\infty$  ( $\mathbf{F} \in \mathbb{R}^{n \times m}$  is the design matrix of the features of the samples of a task), while in inner-Ridge we have  $\lambda_{\max} := \frac{1}{n}\|\mathbf{F}\|^2$ . Note that this means  $\lambda$  is dynamically changing as we train because  $\mathbf{F}$  changes. However we never backpropagate through  $\lambda_{\max}$  (we block the gradient from flowing). Thus, in all figures, we report  $\epsilon = \lambda/\lambda_{\max}$ .

D.2.4. EXPERIMENTS VIOLATING ASSUMPTIONS

In this section, we explore variations of the experiments of Section 5, but this time the assumptions of Theorem 3.1 are violated.

Fig. 8 shows different degrees of violation of Assumption 3.6. We consider the cases where  $\mathcal{S} := \{\{1, 2\}, \{3, 4\}, \{5, 6\}\}$  (block size = 2),  $\mathcal{S} := \{\{1, 2, 3\}, \{4, 5, 6\}\}$  (block size = 3) and  $\mathcal{S} := \{\{1, 2, 3, 4, 5, 6\}\}$  (block size = 6). Note that the latter case corresponds to having no sparsity at all in the ground-truth model, *i.e.*, all tasks require all features. The reader can verify that these three cases indeed violate Assumption 3.6. In all cases, the distribution  $p(\mathcal{S})$  puts uniform mass over its support  $\mathcal{S}$ . Similarly to the experiments from the main text,  $\mathbf{w} := \bar{\mathbf{w}} \odot \mathbf{s}$ , where  $\bar{\mathbf{w}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and  $\mathbf{s} \sim p(\mathcal{S})$  ( $\mathbf{s}$  is the binary representation of the set  $\mathcal{S}$ ). Overall, we can see that inner-Lasso does not perform as well when Assumption 3.6 is violated. For example, when there is no sparsity at all (block size = 6), inner-Lasso performs poorly and is even surpassed by inner-Ridge. Nevertheless, for mild violations (block size = 2), disentanglement (as measured by MCC) remains reasonably high. We further notice that all methods obtain very good R score in all settings. This is expected in light of Theorem B.4, which guarantees identifiability up to linear transformation without requiring Assumption 3.6.

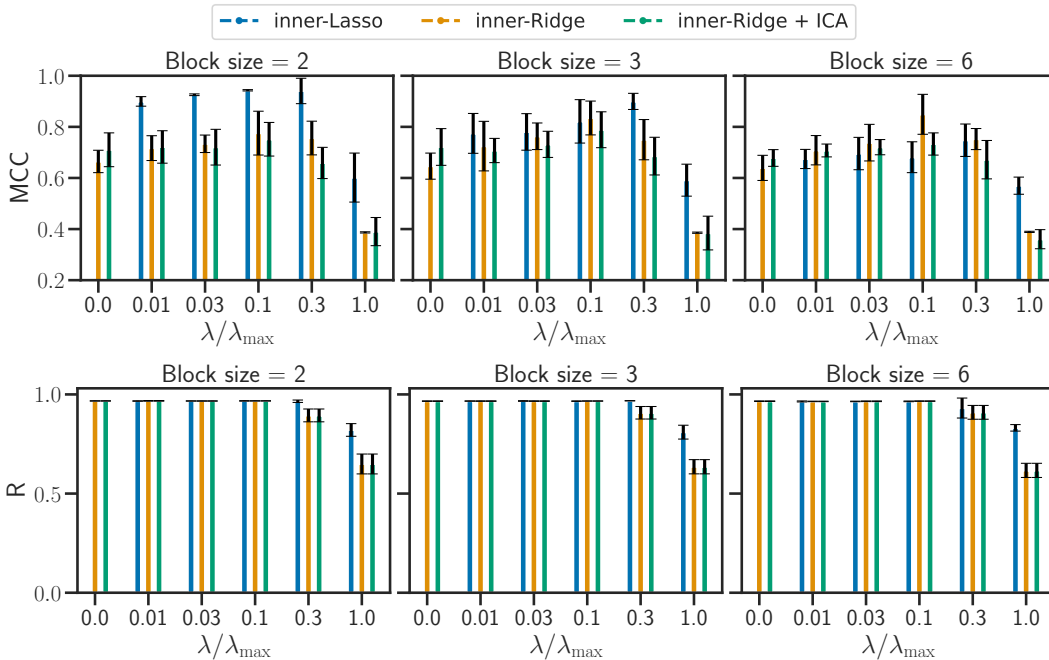


Figure 8. Disentanglement (MCC, top) and prediction (R Score, bottom) performances for inner-Lasso, inner-Ridge and inner-Ridge combined with ICA as a function of the regularization parameter. The metrics are plotted for multiple value of block size for the support. Block size = 6 corresponds to no sparsity in the ground truth coefficients.

Fig. 9 presents experiments that are identical to those of Fig. 4 in the main text, except for how  $\mathbf{w}$  is generated. Here, the components of  $\mathbf{w}$  are sampled independently according to  $w_i \sim \text{Laplace}(\mu = 0, b = 1)$ . We note that, under this process, the probability that  $w_i = 0$  is zero. This means all features are useful and Assumption 3.6 is violated. That being said, due to the fat tail behavior of the Laplacian distribution, many components of  $\mathbf{w}$  will be close to zero (relatively to its variance). Thus, this can be thought of as a weaker form of sparsity where many features are relatively unimportant. Fig. 9 shows that

inner-Lasso can still disentangle very well. In fact, the performance is very similar to the experiments that presented actual sparsity (Fig. 4).

#### D.2.5. EXPERIMENTS WITH REGULARIZATION IN THE OUTER PROBLEM

Theorem B.6 presented an alternative optimization problem to that of Theorem 3.1 to learn a disentangled representation. Appendix B.3 presented a tractable relaxation of this alternative. The essential operational difference is that the sparsity regularization appears in the outer problem instead of the inner problem. Figure 10 shows this alternative works as well empirically. Details in the caption.

#### D.2.6. VISUAL EVALUATION

Fig. 11 to 14 show how various learned representations respond to changing a single factor of variation in the image (Higgins et al., 2017, Figure 7.A.B). We see what was expected: the higher the MCC, the more disentangled the learned features appear, thus validating MCC as a good metric for disentanglement. See captions for details.

#### D.2.7. ADDITIONAL METRICS FOR DISENTANGLEMENT

We implemented metrics from the DCI framework (Eastwood & Williams, 2018) to evaluate disentanglement. 1) DCI-Disentanglement: How many ground truth latent components are related to a particular component of the learned latent representation; 2) DCI-Completeness: How many learned latent components are related to a particular component of the ground truth latent representation. Note that for the definition of disentanglement used in the present work Definition 1.1, we want both DCI-disentanglement and DCI-completeness to be high.

The DCI framework requires a matrix of relative importance. In our implementation, this matrix is the coefficient matrix resulting from performing linear regression with inputs as the learned latent representation  $\mathbf{f}_{\hat{\theta}}(\mathbf{x})$  and targets as the ground truth latent representation  $\mathbf{f}_{\theta}(\mathbf{x})$ , and denote the solution as the matrix  $W$ . Further, denote by  $I = |W|$  as the importance matrix, as  $I_{i,j}$  denotes the relevance of inferred latent  $\mathbf{f}_{\hat{\theta}}(\mathbf{x})_j$  for predicting the true latent  $\mathbf{f}_{\theta}(\mathbf{x})_i$ .

Now, for computing DCI-disentanglement, we normalize each row of the importance matrix  $I[i, :]$  by its sum so that it represents a probability distribution. Then disentanglement is given by  $\frac{1}{m} \times \sum_i^m 1 - H(I[i, :])$ , where  $H$  denotes the entropy of a distribution. Note that for the desired case of each ground truth latent component being explained by a single inferred latent component, we would have  $H(I[i, :]) = 0$  as we have a one-hot vector for the probability distribution. Similarly, for the case of each ground truth latent component being explained uniformly by all the inferred latents,  $H(I[i, :])$  would be maximized and hence the DCI score would be minimized. To compute the DCI-completeness, we first normalize each column of the importance matrix  $I[:, j]$  by its sum so that it represents a probability distribution and then compute  $\frac{1}{m} \times \sum_i^m 1 - H(I[:, j])$ .

Figure 15 shows the results for the 3D Shapes experiments (Section 5) with the DCI metric to evaluate disentanglement. Notice that we find the same trend as we had with the MCC metric 4, that inner-Lasso is more robust to correlation between the latent variables, and inner-Ridge + ICA performance drops down significantly with increasing correlation.

### D.3. Meta-learning experiments

**Experimental settings.** We evaluate the performance of our meta-learning algorithm based on a group-sparse SVM learners on the *miniImageNet* (Vinyals et al., 2016) dataset. Following the standard nomenclature in few-shot classification (Hospedales et al., 2021) with  $k$ -shot  $N$ -way, where  $N$  is the number of classes in each classification task, and  $k$  is the number of samples per class in the training dataset  $\mathcal{D}_t^{\text{train}}$ , we consider the experimental setting 5-shot 5-way. We use the same residual network architecture as in (Lee et al., 2019), with 12 layers and a representation of size  $p = 1.6 \times 10^4$ .

**Technical details.** The objective of Problem (10) is composed of a smooth term and block separable non-smooth term, hence it can be solved efficiently using proximal block coordinate descent (Tseng, 2001). Although Theorem 3.1 is not directly applicable to the meta-learning formulation proposed in this section, we conjecture that similar techniques could be reused to prove an identifiability result in this setting. As in Section 3.4, the argmin differentiation of the solution of Problem (10) can be done using implicit differentiation (Bertrand et al., 2022).

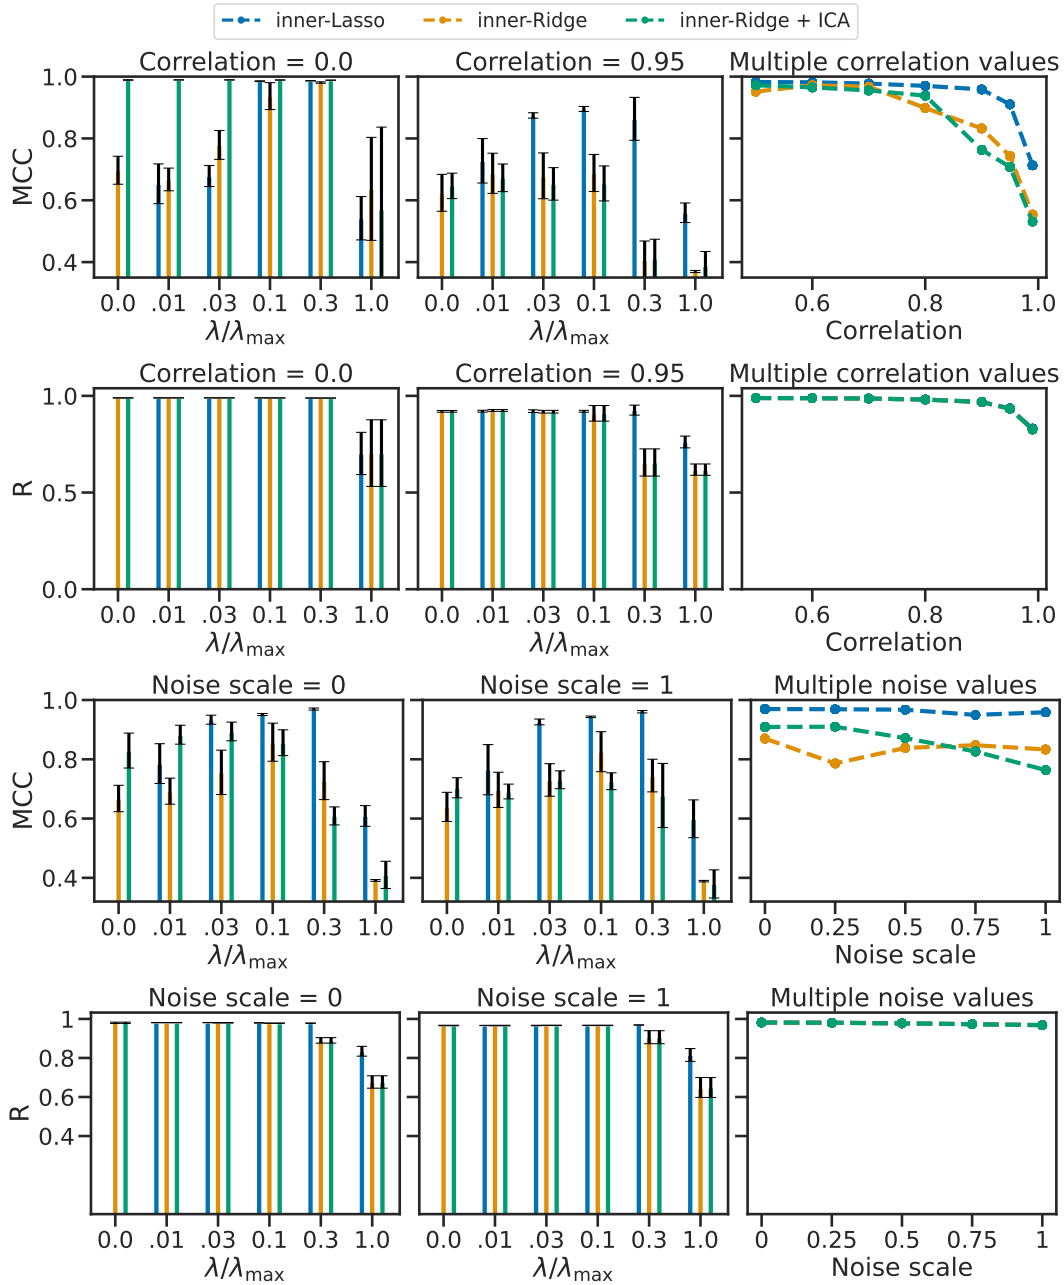


Figure 9. Same experiment as Fig. 4, but the task coefficient vectors  $w$  are sampled from a Laplacian distribution (instead of what was described in Appendix D.2.1). Performance is barely affected, showing some amount of robustness to violations of Assumption 3.6.

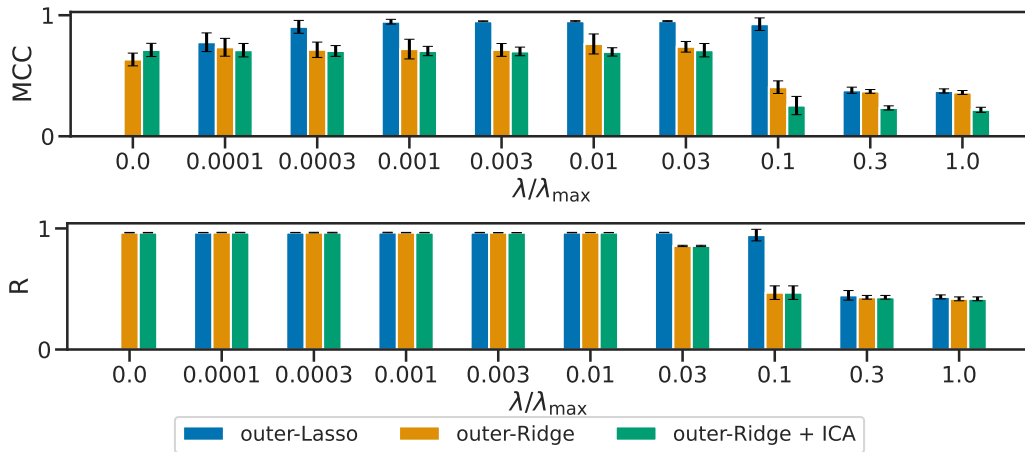


Figure 10. **outer-Lasso** solves Problem (62) (with regularization in the outer problem) while **outer-Ridge** solves the same problem but with an  $L_2$ -norm instead of  $L_{2,1}$ . The method **outer-Ridge + ICA** is outer-Ridge with an additional step of linear ICA on top of the learned representation. The results obtained are very similar to the main results of Fig. 4 and 6. In this dataset, the latents are sampled from  $p_{\alpha=1, \rho=0.9}(\mathbf{z})$  (See Appendix D.2.1) and the weight coefficients are sampled from the binomial-Gaussian process described in Appendix D.2.1.

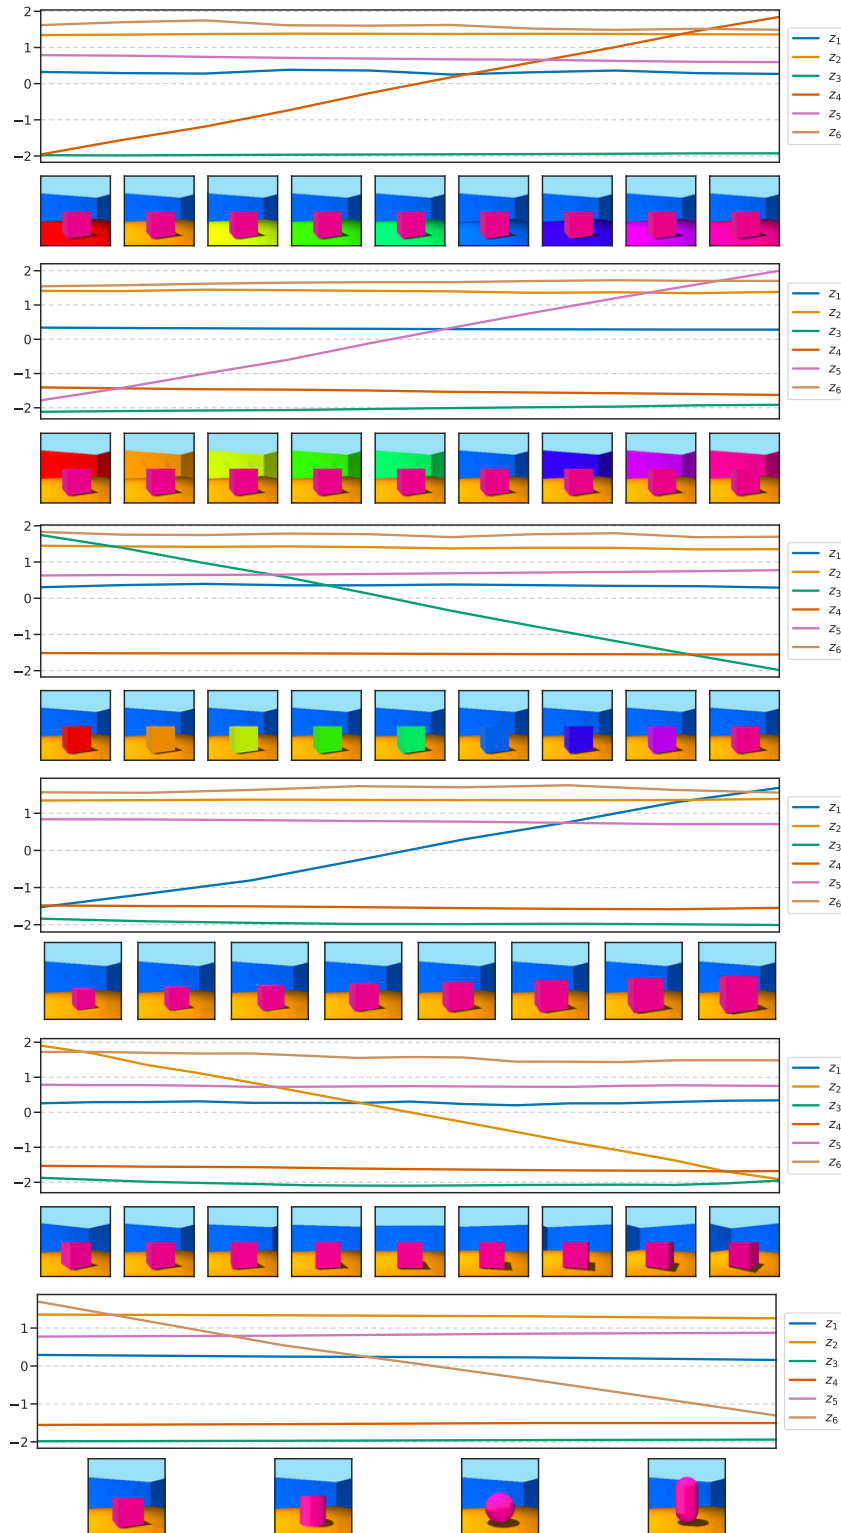


Figure 11. Varying one factor at a time in the image and showing how the learned representation varies in response. This representation was learned by **inner-Lasso** (best hyperparameter) on a dataset with **0 correlation between latents** and a noise scale of 1. The corresponding **MCC is 0.99**. We can see that varying a single factor in the image always result in changing a single factor in the learned representation.



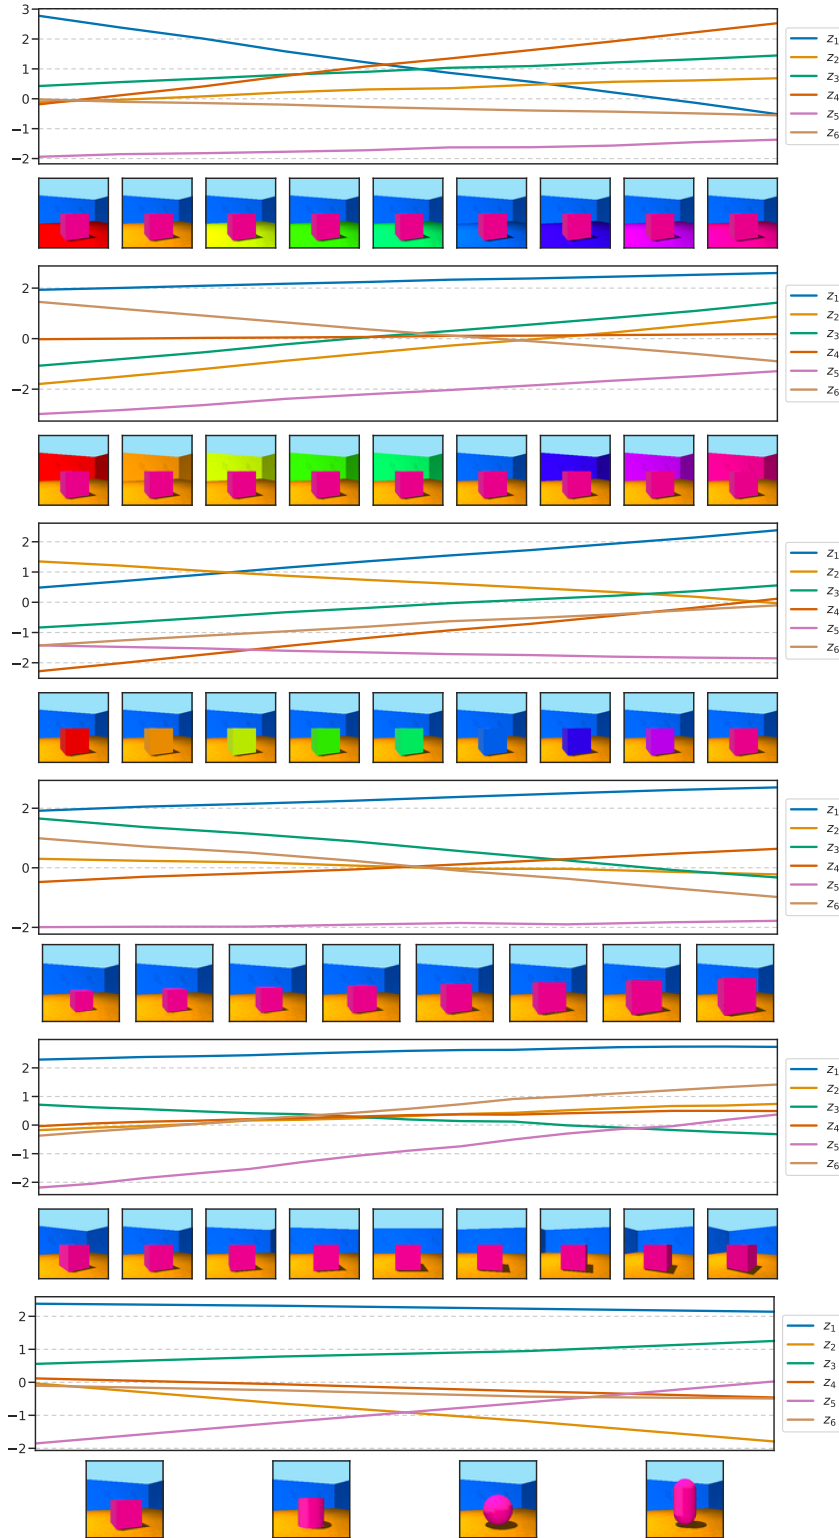


Figure 12. Varying one factor at a time in the image and showing how the learned representation varies in response. This representation was learned **without regularization** of any kind (*i.e.*, with inner-Ridge with regularization coefficient equal to zero) on a dataset with **0 correlation between** and a noise scale of 1. The corresponding **MCC is 0.63**. We can see that varying a single factor in the image result in changing multiple factors in the learned representation, *i.e.*, the representation is not disentangled.

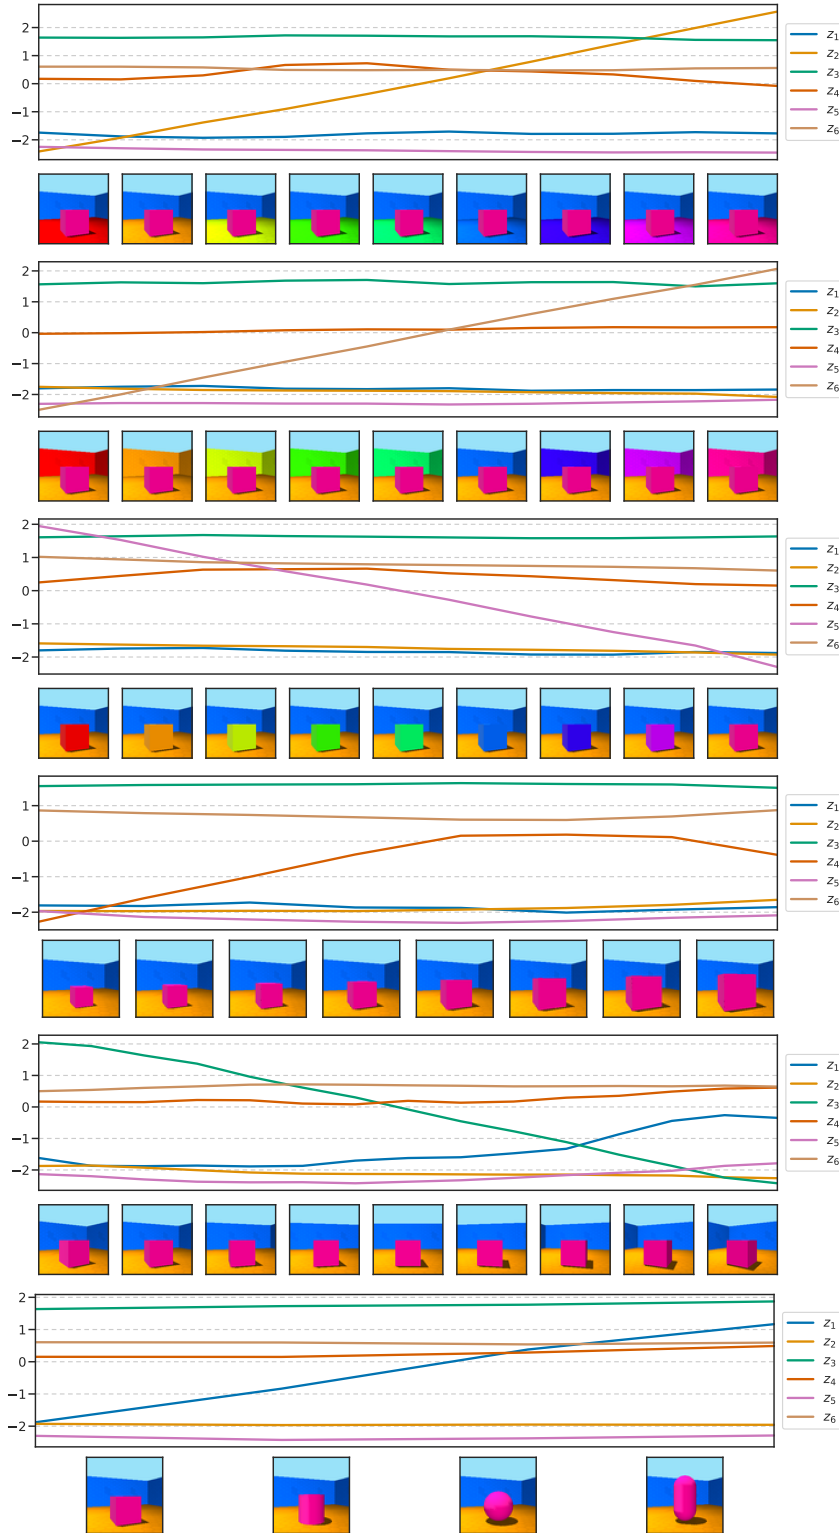


Figure 13. Varying one factor at a time in the image and showing how the learned representation varies in response. This representation was learned with **inner-Lasso** (best hyperparameter) on a dataset with **correlation 0.9 between latents** and a noise scale of 1. The corresponding **MCC is 0.96**. Qualitatively, the representation appears to be well disentangled, but not as well as in Fig. 11 (reflected by a drop in MCC of 0.03).

## Synergies Between Disentanglement and Sparsity

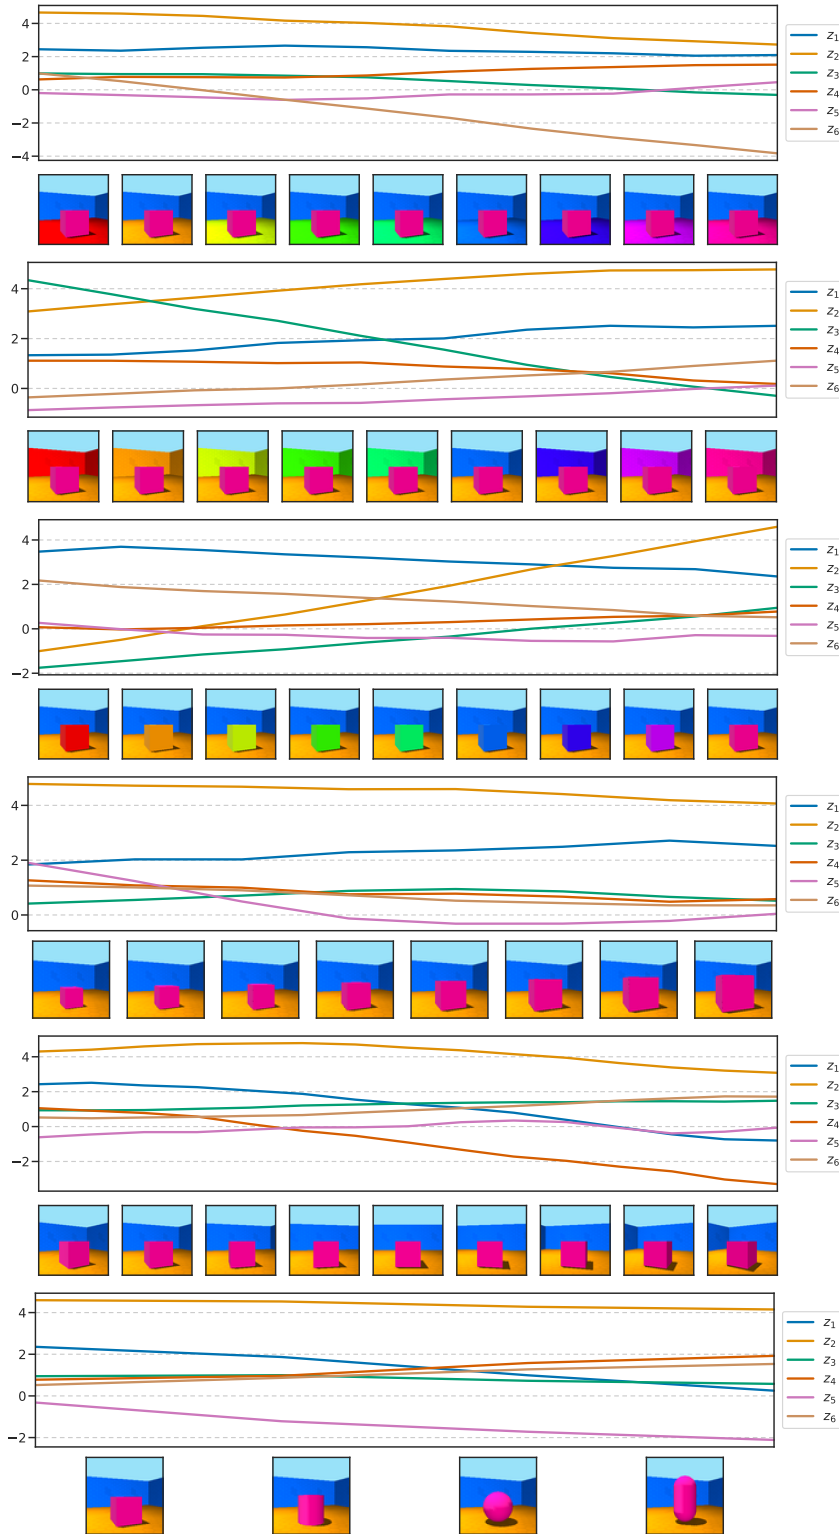


Figure 14. Varying one factor at a time in the image and showing how the learned representation varies in response. This representation was learned with **inner-Ridge** (best hyperparameter) on a dataset with **correlation 0.9 between latents** and a noise scale of 1. The corresponding **MCC is 0.79**. For most latent factors, we cannot identify a dominant feature, except maybe for background and object colors. The representation appears more disentangled than Fig. 12, but less disentangled than Fig. 13, as reflected by their corresponding MCC values.

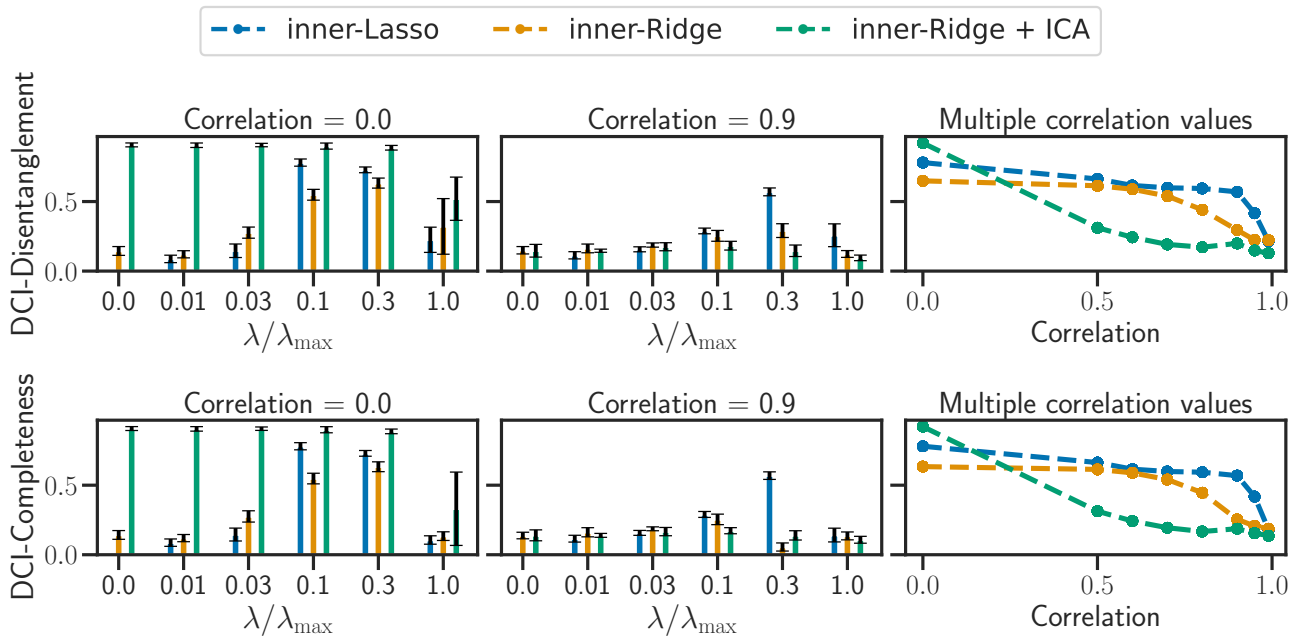


Figure 15. Disentanglement performance (DCI) for inner-Lasso, inner-Ridge and inner-Ridge combined with ICA as a function of the regularization parameter (left and middle). The right columns shows performance of the best hyperparameter for different values of correlation and noise. The top row shows the results for the disentanglement metric of DCI and the bottom row shows the results for the completeness metric of DCI.