

Revisiting Disentanglement in Downstream Tasks: A Study on Its Necessity for Abstract Visual Reasoning

Ruiqian Nai^{1, 2, 3}, Zixin Wen⁴, Ji Li¹, Yuanzhi Li⁴, Yang Gao^{1, 2, 3*}

¹Tsinghua University, Beijing, China

²Shanghai Artificial Intelligence Laboratory, Shanghai, China

³Shanghai Qi Zhi Institute, Shanghai, China

⁴Carnegie Mellon University, Pittsburgh, PA, USA

{nrq22, ji-li20}@mails.tsinghua.edu.cn, {gaoyangiiis}@tsinghua.edu.cn, {zixinw, yuanzhil}@andrew.cmu.edu

Abstract

In representation learning, a disentangled representation is highly desirable as it encodes generative factors of data in a separable and compact pattern. Researchers have advocated leveraging disentangled representations to complete downstream tasks with encouraging empirical evidence. This paper further investigates the necessity of disentangled representation in downstream applications. Specifically, we show that dimension-wise disentangled representations are unnecessary on a fundamental downstream task, abstract visual reasoning. We provide extensive empirical evidence against the necessity of disentanglement, covering multiple datasets, representation learning methods, and downstream network architectures. Furthermore, our findings suggest that the informativeness of representations is a better indicator of downstream performance than disentanglement. Finally, the positive correlation between informativeness and disentanglement explains the claimed usefulness of disentangled representations in previous works. The source code is available at <https://github.com/Richard-coder-Nai/disentanglement-lib-necessity>.

1 Introduction

Disentanglement has been considered an essential property of representation learning (Bengio, Courville, and Vincent 2013; Peters, Janzing, and Schölkopf 2017; Goodfellow, Bengio, and Courville 2016; Bengio, LeCun et al. 2007; Schmidhuber 1992; Lake et al. 2017; Tschannen, Bachem, and Lucic 2018). Disentanglement is defined as a dimension-wise relationship, wherein a representation dimension should capture information from exactly one factor and vice versa (Locatello et al. 2019b; Higgins et al. 2016; Kim and Mnih 2018; Chen et al. 2018; Eastwood and Williams 2018; Ridgeway and Mozer 2018; Kumar, Sattigeri, and Balakrishnan 2017; Do and Tran 2019). Such property is analogous to biological mechanisms, as neurons in brains are specialized for specific tasks, with some aligned to axes of data generative factors (Higgins et al. 2021; Whittington et al. 2022). Moreover, disentangled representations offer a compact and separable structure. Therefore, they are believed to aid in compositional generalization and reasoning, potentially leading to improved performance in downstream tasks (Bengio, Courville, and Vincent 2013).

These purported advantages have been verified on multiple downstream tasks. For instance, abstract visual reasoning (van Steenkiste et al. 2019), fairness (Locatello et al. 2019a), and out-of-distribution (OOD) generalization (Dittadi et al. 2020). Disentangled representations result in better downstream performance, faster learning, and a strong correlation with task success. Based on these promising empirical results, adopting disentangled representations is considered a wise decision for completing downstream tasks.

However, Cao et al. (2022) demonstrates that contrastive learning, does not produce fully disentangled representations. Yet contrastive pre-training methods achieve remarkable downstream performances (He et al. 2020; Caron et al. 2021; He et al. 2022; Oquab et al. 2023). These findings may contradict earlier claims regarding the benefits of developing disentangled representations, as such disentanglement does not appear to play a significant role in the success of downstream tasks. This potential conflict with the previously claimed usefulness of disentanglement motivates us to reassess the role of disentanglement for downstream tasks. Are disentangled representations *necessary* for downstream tasks? If not, how do we explain the previously reported benefits?

We choose abstract visual reasoning as the testbed to investigate the necessity of disentanglement. In this task, intelligent agents are asked to take human IQ tests, i.e., predict the missing panel of Raven’s Progressive Matrices (RPMs) (Raven 1941). To solve this task, it is essential to comprehend and employ the generative factors of data, as they serve as reasoning attributes. The task is completed in a **two-staged** fashion: (1) extracting RPMs’ representations in an unsupervised manner, (2) then performing abstract reasoning training based on learned representations. Therefore, it is an essential and popular benchmark for disentangled representation learning (van Steenkiste et al. 2019; Locatello et al. 2020; Schölkopf et al. 2021). In practice, among various researched downstream tasks (van Steenkiste et al. 2019; Locatello et al. 2019a; Dittadi et al. 2020), disentanglement has been shown to improve sample efficiency and final performance on abstract reasoning tasks (van Steenkiste et al. 2019). Researchers have therefore suggested the use of disentangled representations (Steenbrugge et al. 2018; van Steenkiste et al. 2019; Małkiński and Mańdziuk 2022).

We conduct an extensive empirical study that uses abstract

*Corresponding author.

reasoning tasks to investigate the role of disentanglement for downstream tasks. We train 720 representation learning models on two datasets, including disentanglement and general-purpose methods. We then train 5 WReNs (Barrett et al. 2018) and 5 Transformers (Vaswani et al. 2017; Hahne et al. 2019) using the outputs of each representation learning model to perform abstract reasoning, yielding a total of 7200 abstract reasoning models. Our contributions are as follows.

- We conduct a comprehensive exploration into the impact of disentanglement on downstream tasks, by introducing both disentangled and general-purpose representations and employing multiple methods (WReN and Transformer) to complete the downstream task.
- We show the unnecessary by highlighting the significance of *informativeness* over disentanglement. Informativeness measures *what* information the representation has learned (Eastwood and Williams 2018). Previous studies’ analysis on informativeness is scarce and thus overlook its importance (van Steenkiste et al. 2019; Locatello et al. 2019a; Dittadi et al. 2020; Trauble et al. 2021).
- We show that informativeness is the underlying factor behind the previously argued usefulness of disentanglement, as we observe limited extra benefits of disentanglement when informativeness is closely matched.

2 Related Work

Disentangled representation learning. In this paper, we adopt the widely accepted definition of disentanglement: a one-to-one mapping between representation dimensions and generative factors of data, which we term “dimension-wise disentanglement”. It requires that each representation dimension encode only one factor and vice versa (Locatello et al. 2019b; Eastwood and Williams 2018; Kumar, Sattigeri, and Balakrishnan 2017; Do and Tran 2019). While some studies suggest relaxing this constraint, introducing new properties, or expanding the definition of disentanglement, (Higgins et al. 2018; Wang and Jordan 2021; Roth et al. 2022; Eastwood et al. 2022), our focus remains on this well-established definition. Another line of work related to disentangled representation learning is the Independent Component Analysis (ICA) (Comon 1994). ICA aims to recover independent components of the data.

Based on the dimension-wise definition, researchers develop methods and metrics. SOTA disentanglement methods are mainly variants of generative methods (Higgins et al. 2016; Kim and Mnih 2018; Burgess et al. 2018; Kumar, Sattigeri, and Balakrishnan 2017; Chen et al. 2018, 2016; Jeon, Lee, and Kim 2018; Lin et al. 2020; Leeb et al. 2020). Corresponding metrics are designed (Higgins et al. 2016; Kim and Mnih 2018; Chen et al. 2018; Eastwood and Williams 2018; Kumar, Sattigeri, and Balakrishnan 2017; Cao et al. 2022).

Downstream tasks. Several works conduct empirical studies on downstream tasks to support the believed benefits of disentanglement (Bengio, Courville, and Vincent 2013; Do and Tran 2019), including abstract reasoning (van Steenkiste et al. 2019), fairness (Locatello et al. 2019a), and OOD generalization (Dittadi et al. 2020). Provided with positive

empirical results, these works advocate using disentanglement to complete downstream tasks. Among these works, van Steenkiste et al. (2019) reported the most encouraging evidence from abstract reasoning tasks. Disentanglement is more significant than other representation properties, especially in limited samples. Therefore, we adopt their settings and investigate the necessity of disentanglement on the same tasks at a few-sample regime. However, there are a few issues with their study. Firstly, it underestimates factors’ linear classification accuracy, yielding a weaker correlation between informativeness and downstream performance (see Figure 8 in Appendix B.1). Moreover, it paid insufficient attention to the analysis of informativeness. We address these issues and show the unnecessary of disentanglement on a broader range of representations and downstream methods.

Trauble et al. (2021) delves into the role of sim2real transfer in reinforcement learning tasks using real robots, and their findings closely align with ours. However, their scope is still confined to variants of VAEs. Our work is complementary to theirs on a more fundamental and straightforward mechanism. Instead of the sophisticated OOD generalization tasks, we focus on the essential setting of abstract visual reasoning.

Abstract visual reasoning has been a popular benchmark to measure the representation’s downstream performance, especially in disentanglement literature (Steenbrugge et al. 2018; van Steenkiste et al. 2019; Dittadi et al. 2020; Locatello et al. 2020; Scholkopf et al. 2021). The most common type is the Raven’s Progressive Matrices (RPMs) (Raven 1941). To solve RPMs, one is asked to complete the missing panel of a 3×3 grid by exploring the logical relationships of 8 context panels. Moreover, abstract visual reasoning is a well-developed benchmark for representation learning. Given that it is coupled with a principle treatment of generalization (Fleuret et al. 2011), a neural network can not solve reasoning tasks by simply memorizing superficial statistical features.

3 Downstream Benchmark: Abstract Visual Reasoning

In this section, we’ll introduce the abstract visual reasoning task and outline the downstream benchmark framework, including representation learning methods, metrics, and abstract reasoning models.




Figure 1: An example of RPM on *3DShapes* from van Steenkiste et al. (2019).

3.1 Abstract Visual Reasoning as a Two-staged Task

The abstract visual reasoning tasks are highly inspired by the famous human IQ test, Raven’s Progressive Matrices

(RPMs) (Raven 1941). Figure 1 shows an RPM question in our evaluation dataset. There are eight context panels and one missing panel in the left part of the figure. The context panels are arranged following some logical rules across rows. During the test, the subject must pick one of the six candidates in the right part to fix the missing panel. The goal is to maintain the logical relationships given by the contexts. More details of RPMs are available in Appendix A.4.

We adopt RPMs as a downstream benchmark following van Steenkiste et al. (2019). To measure the necessity of disentanglement for downstream tasks, we separate the evaluation process into two stages: **(1)** In Stage-1, representation learning models extract representations from images of which RPMs consist, and **(2)** in Stage-2, abstract reasoning models predict the missing panels from the frozen representations of contexts and answer candidates. Correspondingly, we denote representation learning models as **Stage-1 models** while abstract reasoning models as **Stage-2 models**. For Stage-1, we measure the representation properties, including disentanglement and informativeness. A diverse set of Stage-1 and Stage-2 models are trained, yielding multiple samples from the joint distribution of representation metric scores and downstream accuracy. Finally, we study the relationships between representation qualities and downstream performance. We aim to investigate whether more disentangled representations perform better on abstract reasoning tasks.

3.2 Background of Representation Learning

Disentangled representation learning methods. The seminal works of Higgins et al. (2016) and Chen et al. (2016) embody disentanglement by augmenting deep generative models (Kingma and Welling 2013; Goodfellow et al. 2014). For disentangled representation learning methods, following van Steenkiste et al. (2019), we focus on a family of VAEs with disentanglement inductive bias denoted as DisVAEs. DisVAEs’ objective summarizes augmentations of SOTA methods. Namely, β -VAE (Higgins et al. 2016), Annealed-VAE (Burgess et al. 2018), β -TCVAE (Chen et al. 2018), FactorVAE (Kim and Mnih 2018), and DIP-VAE (Kumar, Sattigeri, and Balakrishnan 2017). They achieve disentanglement mainly by encouraging independence between representation dimensions. Please refer to Appendix A.2 for details.

General-purpose representation learning methods. In our study, methods not (explicitly) encouraging disentanglement are called general-purpose methods. We take a set of BYOL (Grill et al. 2020) with the same size as DisVAEs as representatives. BYOL is a negative-free contrastive learning method. It creates different “views” of an image by data augmentation and pulls together their distance in representation space. To avoid collapsing to trivial representations, a predictor appending to one of the siamese encoders and exponential moving average update strategy (He et al. 2020) are employed. It does not encourage disentanglement due to the lack of regularizers. Indeed, the empirical evidence in Cao et al. (2022) demonstrates that representations learned by BYOL have weak disentanglement properties.

Representation property metrics. Considered properties of representations cover two axes of metrics: disentanglement

metrics and informativeness metrics (Eastwood and Williams 2018; Eastwood et al. 2022). We include *BetaVAE* score (Higgins et al. 2016), *FactorVAE* score (Kim and Mnih 2018), *Mutual Information Gap* (Chen et al. 2018), *SAP* (Kumar, Sattigeri, and Balakrishnan 2017), and *DCI Disentanglement* (Eastwood and Williams 2018). Locatello et al. (2019b) proves their agreement on VAE methods with extensive experiments. Though their measurements are different, their results are positively correlated. To facilitate comparison across different representation sizes, we also include the *MED* metric (Cao et al. 2022). On the other hand, informativeness requires representations to encode enough information about factors (Eastwood and Williams 2018). We employ *Logistic Regression* (LR) and *Linear Regression* in this work. They use a linear model to classify or regress the values of generative factors. Given the weak capacity of linear models, a higher LR accuracy or lower regression error ensures that sufficient information is explicitly encoded. However, it does not emphasize a dimension-wise encoding pattern like disentanglement. To distinguish, we term the property indicated by LR and linear regression as *informativeness* (Eastwood and Williams 2018).

3.3 Background of Methods for Abstract Reasoning

In Stage-1, we extract representations of eight context panels (the left part of Figure 1) and six answer candidates (the right part of Figure 1). Then in Stage-2, downstream models perform abstract reasoning from the (frozen) representations. Abstract reasoning models evaluate whether filling the blank panel by a candidate follows the logical rules given by contexts. For a trial T_i of one candidate $a_i \in A = \{a_1, \dots, a_6\}$ and eight context panels $C = \{c_1, \dots, c_8\}$, its score is calculated as follows:

$$Y_i = \text{Stage2}(\text{Stage1}(T_i)),$$

$$\text{Stage1}(T_i) = \{\text{Stage1}(c_1), \dots, \text{Stage1}(c_8), \text{Stage1}(a_i)\},$$

where Y_i is the score of trial T_i , $\text{Stage1}(\cdot)$, $\text{Stage2}(\cdot)$ denote the forward process of the Stage-1 and Stage-2 models, and $\text{Stage1}(T_i)$ is the representations of contexts and candidate a_i . After evaluating all trials $\{T_1, T_2, \dots, T_6\}$, the output answer is $\arg \max_i Y_i$.

We implement two different structures of Stage-2 models, namely, WReN (Barrett et al. 2018) and Transformer (Vaswani et al. 2017; Hahne et al. 2019). First, we employ an MLP or a Transformer to embed an RPM trial. Then, an MLP head predicts a scalar score from the embeddings.

4 Experiments

In this section, we conduct a systematic empirical study about representation properties’ impacts on downstream performance. First, we introduce our experimental conditions in Section 4.1. Then, in Section 4.2, we demonstrate that disentangled and general-purpose representations have similar performance, confirming the following experiment. Finally, we demonstrate how informativeness proves to be a stronger indicator of downstream performance in Section 4.3.

4.1 Experiments Setup

We build upon the experiment conditions of van Steenkiste et al. (2019). Abstract visual reasoning tasks, i.e., RPMs, are solved through a two-stage process: data $\xrightarrow{\text{Stage-1}}$ representations $\xrightarrow{\text{Stage-2}}$ RPM answers. We first train Stage-1 models in an unsupervised manner and evaluate their disentanglement and informativeness. Then, Stage-2 models are trained and evaluated on downstream tasks, yielding an abstract reasoning accuracy of a representation. Provided with a large amount of (representation property score, downstream performance) pairs, we conduct a systematic study to investigate the necessity of disentanglement. More implementation details are available in Appendix A.

Datasets. We replicate the RPM generation protocol in van Steenkiste et al. (2019). The panel images consist of disentanglement benchmark image datasets, namely, *Abstract dSprites* (Matthey et al. 2017; van Steenkiste et al. 2019) and *3DShapes* (Burgess and Kim 2018). The rows of RPMs are arranged following the logical *AND* of ground truth factors. As for hardness, we only reserve *hard-mixed*, whose contexts and candidates are more confusing. According to the generation process, the size of generated RPMs is sufficiently large (about 10^{144}), allowing us to produce fresh samples throughout training.

Reference models. Stage-1 models include 360 disentangled VAEs (denoted as DisVAEs) and 360 BYOLs, covering both disentangled and general-purpose representation learning methods. A diverse set of configurations are included. According to the histograms in Appendix C.4, our choices of Stage-1 models span various disentanglement and informativeness scores. For Stage-2, we train 10 Stage-2 models (5 WReNs and 5 Transformers) for every Stage-1 model. These configurations are randomly sampled from a search space described in Appendix A.3 and shared across Stage-1 models to ensure fair comparisons.

Training protocol. Training is conducted two-staged. Firstly, we train Stage-1 models in an unsupervised manner on the dataset consisting of RPMs’ panels, i.e., *Abstract dSprites* or *3DShapes*. For DisVAE models, we use the training protocol of van Steenkiste et al. (2019), while for BYOL models, we follow Cao et al. (2022). In Stage-2, all models are trained for 10K iterations with a batch size of 32. After every 100 iterations, we evaluate the accuracy on newly generated 50 mini-batches of unseen RPM samples for validation and another 50 mini-batches for testing.

Evaluation protocol. We first evaluate the two stages separately. Then, we analyze the relationship between the two stages, i.e., representation properties and downstream performance. Specifically, to investigate the necessity of disentanglement, we are interested in whether more disentangled representations lead to better downstream performance. Further, can we find another metric that better accounts for downstream performance? Therefore, for Stage-1, we employ representation metrics described in Section 3.2 to measure two aspects: disentanglement and informativeness. For all Stage-1 models, we compute the following metric scores: *BetaVAE* score, *FactorVAE* score, *MIG*, *SAP*, *MED*, *LR* accu-

racy, and negative normalized mean squared error of *Linear Regression*. *DCI Disentanglement* is only evaluated for DisVAEs. Since we follow the tree-based implementation used in previous studies (Locatello et al. 2019b; van Steenkiste et al. 2019; Dittadi et al. 2020; Träuble et al. 2021) for easy comparison. We also provide *DCI Disentanglement* score based on LASSO on both DisVAEs and BYOL (see Figure 9 and Appendix B.2). For Stage-2, we inspect accuracy on newly generated test sets every 100 iterations, yielding accuracy for multiple training steps. Since every step sees fresh samples, we employ training curves to measure sample efficiency following van Steenkiste et al. (2019); Träuble et al. (2021).

To summarize the downstream performance of a Stage-1 model, over 5 WReNs or 5 Transformers in Stage-2, we report the mean accuracy denoted as $\overline{\text{WReN}}$ or $\overline{\text{Trans}}$. Finally, we calculate the rank correlation (Spearman) between the mean performance of Stage-1 models ($\overline{\text{WReN}}$ and $\overline{\text{Trans}}$) at certain Stage-2 steps and their Stage-1 metric scores. Rank correlation has been widely adopted in the literature studying disentanglement downstream tasks (van Steenkiste et al. 2019; Locatello et al. 2019a; Dittadi et al. 2020; Locatello et al. 2020). A larger correlation indicates a higher significance of the representation property on downstream performance.

4.2 Preliminary Experiments on Representation Variants

Previous studies have primarily focused on DisVAEs (van Steenkiste et al. 2019; Locatello et al. 2019a; Dittadi et al. 2020; Träuble et al. 2021). In an endeavor to ensure the generalizability of our findings, two distinct representation variants are employed: disentangled representations and general-purpose representations. In this section, we present the preliminary experiments conducted to compare the performance of these representations on the downstream task. Establishing comparable performance between these representation variants is important to substantiate the rationale behind considering both representation types in our analysis.

Comparative Performance Analysis. We show the downstream performance of different families of learning models described in Section 4.1, including disentanglement-oriented (DisVAEs) and general-purpose (DisVAEs) representations. To ensure equitable comparison, we select the most effective representations from a parameter search space of equal size (360 DisVAEs and 360 BYOLs). Each representation from Stage-1 was used to train 5 WReNs and 5 Transformers, and their performance was measured by averaging their downstream accuracy ($\overline{\text{WReN}}$, $\overline{\text{Trans}}$). In this section, we present the models demonstrating the highest $\overline{\text{WReN}}$ and $\overline{\text{Trans}}$ performance.

We first determine the final accuracy of disentangled representations versus less disentangled ones. In Table 1, we present the *best* results for $\overline{\text{WReN}}$ and $\overline{\text{Trans}}$ achieved across various datasets. We select checkpoints to evaluate based on validation accuracy. It suggests that BYOL performs slightly better than DisVAEs. In particular, the best $\overline{\text{WReN}}$ and $\overline{\text{Trans}}$ of BYOL are higher than that of DisVAEs’. However, the MED scores of BYOL checkpoints in Table 1 are significantly




Figure 2: Average test accuracy on *3DShapes* throughout the training. The shaded area indicates the maximum and minimum values. We select the Stage-1 models with best $\overline{\text{WReN}}$ or $\overline{\text{Trans.}}$ among 3600 checkpoints on *3DShapes*. Stage-1 models with disentanglement inductive bias (DisVAEs) are not necessarily better than those without such bias (BYOL) regarding sample efficiency and final accuracy.

Dataset	Stage1	$\overline{\text{WReN}}$	$\overline{\text{Trans.}}$	MED ($\overline{\text{WReN}}$)	MED ($\overline{\text{Trans.}}$)
<i>3DShapes</i>	DisVAEs	84.8(0.91)	87.0(6.36)	0.60	0.68
	BYOL	87.1(4.68)	88.0(2.62)	0.13	0.12
<i>Abstract dSprites</i>	DisVAEs	68.7(1.39)	66.4(7.06)	0.29	0.29
	BYOL	72.2(3.11)	78.1(1.75)	0.13	0.13

Table 1: Downstream test performance (%) and MED scores of different Stage-1 models. For each measurement ($\overline{\text{WReN}}$ and $\overline{\text{Trans.}}$), we report the best among all Stage-1 models. The step with the highest validation accuracy is reported. The numbers in the parentheses are STDs of the 5 scores used in computing $\overline{\text{WReN}}$ or $\overline{\text{Trans.}}$. The MED scores correspond to the checkpoints that achieve the reported $\overline{\text{WReN}}$ and $\overline{\text{Trans.}}$ performances.

lower than those of DisVAEs, indicating that they are less disentangled. Therefore, both entangled and disentangled representations can achieve similar performance, which suggests that it is valuable to investigate the necessity of disentanglement in general purpose and disentanglement methods.

Now we examine another purported benefit of disentanglement: sample efficiency. According to the setting of Stage-2, every step sees fresh samples during training. Therefore, the number of training steps is a fraction of the total number of samples. To gauge sample efficiency, we can observe how accuracy improves over training steps using training curves. This method was used by van Steenkiste et al. (2019) and Träuble et al. (2021). Figure 2 shows overviews of training trajectories of Stage-1 models with the highest performing $\overline{\text{WReN}}$ and $\overline{\text{Trans.}}$ on *3DShapes*. For $\overline{\text{WReN}}$ as Stage-2 models (Figure 2a), BYOL leads at the beginning, then DisVAEs catch up. Finally, BYOL converges at a higher accuracy. In contrast, when Stage-2 models are Transformers, BYOL’s curve grows faster, but DisVAEs and BYOL converge with comparable performance. Generally, the two curves follow nearly identical patterns with small gaps, indicating no clear superiority in terms of sample efficiency between DisVAEs and BYOL. Consequently, disentanglement’s enhancement of sample efficiency can be achieved with less disentangled representations. The same study is conducted on another set of general-purpose models, SimSiam (Chen and He 2021), where we reach similar results (see Appendix C.2).

Ensuring Representation Diversity and Equivalency.

Disentangled and general-purpose representation sets are of uniform size and have a broad spectrum of performance characteristics. Figure 12 displays the metric scores, and Fig-

ure 13 and 14 display the downstream performance of our selected DisVAEs and BYOL. We can see that both representations cover a wide range of metric scores and downstream accuracy (encompassing both superior and inferior ones).

Summary: The preliminary experiments reveal a comparable performance between disentangled and general-purpose representations on the abstract reasoning task. Moreover, both representation categories are of equal size and encapsulate a diverse performance range. This parity validates the subsequent evaluation of the correlation between disentanglement and downstream task performance across these representation variants.

4.3 Informativeness as a Stronger Correlate Than Disentanglement

To investigate the necessity of disentanglement, we analyze how various representation properties, including informativeness and disentanglement metrics, impact downstream performance. We first show that informativeness correlates most with downstream performance, indicating that disentanglement is not necessary. Further, we demonstrate that the previously claimed benefits of disentanglement (Bengio, Courville, and Vincent 2013; Higgins et al. 2016; van Steenkiste et al. 2019; Locatello et al. 2019a; Dittadi et al. 2020), are actually derived from its positive correlation with informativeness.

Recall that we train 720 Stage-1 and 7200 Stage-2 models (see Section 4.1). By taking $\overline{\text{WReN}}$ and $\overline{\text{Trans.}}$ as measurements (average reasoning accuracy over 5 $\overline{\text{WReN}}$ s or 5 Transformers), we yield 720 representations paired with their downstream performance. This section studies the impor-




Figure 3: Rank correlations between $\overline{\text{WReN}}$ or $\overline{\text{Trans.}}$ and representation metrics on *3DShapes*. We denote the step with the highest validation accuracy as “Best”. The brighter the panel, the more correlated the representation metric is with the downstream performance.




Figure 4: Representation metrics versus $\overline{\text{WReN}}$ at step 10000, where Stage-1 models are BYOL, and the dataset is *3DShapes*. We can observe a strong positive correlation between the informativeness metric scores and downstream accuracy.

tance of disentanglement and other representation properties. We must emphasize general trends and overall relationships to show different properties’ significance. Therefore, we follow previous studies (Locatello et al. 2019b; van Steenkiste et al. 2019; Locatello et al. 2019a; Dittadi et al. 2020; Locatello et al. 2020; Träuble et al. 2021) to analyze rank correlation (Spearman) between representation metric scores and downstream performance. If the correlation score is high, we can conclude that the representation property measured by the considered metric score is significant to downstream performance.

Our research builds upon previous studies (Locatello et al. 2019b; van Steenkiste et al. 2019; Locatello et al. 2019a; Dittadi et al. 2020; Locatello et al. 2020; Träuble et al. 2021) by addressing two key issues. Firstly, we have corrected the previous underestimation of informativeness. Secondly, we have expanded the scope of both Stage-1 and Stage-2 models. Figure 8 illustrates how previous works yielded lower informativeness and resulted in lower correlations, leading to an overestimation of the importance of disentanglement. Additionally, the limited representation learning methods only confined to DisVAEs and reasoning models constrained to WReNs have restricted the generalizability of their results.

By addressing these issues, our research leads to contrasting conclusions.

The more significant representation property. We calculate the rank correlation between downstream accuracy with disentanglement and informativeness scores. Meanwhile, we report rank correlation at steps 1K, 2K, 5K, and 10K, and the step with the highest validation accuracy. Because each step involves new samples, the training steps represent a fraction of the sample size. By examining correlations at various training steps, we can determine the impact of representation properties on sample efficiency.

Figure 3 displays rank correlations between representation metric scores and abstract reasoning test accuracy on *3DShapes*. Firstly, at each step throughout the training process, there is always at least one informative metric, such as *Logistic Regression* accuracy (LR) or negative error of linear regressions, that exhibits the most significant correlation with downstream performance. Specifically, LR always correlates more positively than disentanglement metrics. The strong correlation is exploited for all considered models at multiple steps. Since LR and linear regression require sufficient information to be captured and extracted easily from representations, we can conclude that the informativeness

BetaVAE Score (D1)	100	79	58	29	64	61	67	64	Adj.BetaVAE	9	10	5	1	1
FactorVAE Score (D2)	79	100	49	44	71	53	65	53	Adj.FactorVAE	26	19	10	10	10
DCI Disentanglement (D3)	58	49	100	65	58	97	25	26	Adj.DCI-D	19	8	-4	-8	-9
MIG (D4)	29	44	65	100	70	64	6	-4	Adj.MIG	28	9	-5	-9	-8
SAP (D5)	64	71	58	70	100	61	36	28	Adj.SAP	22	5	-6	-8	-8
MED (D6)	61	53	97	64	61	100	24	24	Adj.MED	17	6	-7	-10	-11
LR (I1)	67	65	25	6	36	24	100	87	Adj.Reggression	-30	-15	-4	-3	-5
Regression (I2)	64	53	26	-4	28	24	87	100						
	D1	D2	D3	D4	D5	D6	I1	I2		1000	2000	5000	10000	Best

(a) Overall rank correlations.

(b) Correlations of adjusted metrics.

Figure 5: (a) Correlations between metrics and (b) correlations between adjusted metrics and downstream accuracy when using DisVAEs-WReN pipeline on *3DShapes*. Disentanglement metrics exhibit positive correlations with informativeness. Yet when conditioned on close informativeness, their adjusted versions show mild correlations.

matters most in broad conditions. In contrast, we observe that the importance of disentanglement varies among Stage-1 model families. Disentangled representation learning models (DisVAEs) exhibit strong positive correlations for several disentanglement metrics (but weaker than at least one of informativeness metrics), such as *FactorVAE* score and *DCI Disentanglement*. However, their significance does not apply to BYOL, where the correlation of disentanglement is mild or even negative. In Figure 4, we plot the ($\overline{\text{WReN}}$, metric score) pairs at step 10000. Indeed, for BYOL-WReN on *3DShapes*, we can see the reg-plot provides a good fit of downstream accuracy and informativeness metrics. As for disentanglement metrics, we can see that *BetaVAE* score and *FactorVAE* score suffer from narrow spreads. The regression lines have negative slopes for *MIG*, *SAP*, and *MED*. We conduct a similar analysis on another dataset (*Abstract dSprites*), and another Stage-1 model (SimSiam), and take the same observations. Please refer to Appendix C.4 and Appendix C.2. On the fairness downstream task (Locatello et al. 2019a), we also find that informativeness correlates most (see Appendix C.1 and Table 3).

Summary: The importance of informativeness surpasses that of disentanglement, as evidenced by consistent results across various datasets and model structures.

Understanding for the previously claimed success of disentanglement. Previous works (van Steenkiste et al. 2019; Locatello et al. 2019a; Dittadi et al. 2020; Locatello et al. 2020) have reported empirical evidence backing up the advantages of disentangled representations. Consistently, we observe relatively strong correlations with disentanglement metrics, especially when Stage-1 models are DisVAEs in Figure 3. Based on our conclusion on the significance of the informativeness, we study the DisVAE-WReN case. We provide some insights to explain why the disentanglement metrics have a high correlation to downstream performance in some cases.

The correlations between various metrics have been computed and the results are displayed in Figure 5a. Upon analysis of DisVAEs, it was discovered that there is a strong correlation between informativeness and disentanglement. Additionally, informativeness has a significant correlation with both *FactorVAE* score and *BetaVAE* score. This is evident in Figure 3a where these disentanglement metrics have

a strong correlation with downstream performance. However, other disentanglement metrics have only a mild correlation with informativeness and are ineffective for downstream performance. Therefore, it can be concluded that disentanglement metrics cannot truly predict downstream performance, but informativeness can.

To “purify” the effect of disentanglement, a natural question is: If two representations are of close informativeness, is the more disentangled one more helpful for downstream tasks? For this, we employ adjusted metrics in Locatello et al. (2019a): $\text{Adj. Metric} = \text{Metric} - \frac{1}{5} \sum_{i \in N(\text{LR})} \text{Metric}_i$. For a representation and a certain metric (we care more about disentanglement metrics), we denote its original metric score as Metric . Then we find its 5 nearest neighbors in terms of LR, which we write as $N(\text{LR})$. Finally, the difference between the original metric score and the mean score of the nearest neighbors is reported as adjusted metrics. Intuitively, we calculate the relative disentanglement for representations with close LR.

Figure 5b displays correlations between adjusted metrics and downstream performance. We can find that all adjusted disentanglement metrics correlate mildly with downstream performance. From this, we can see that when informativeness is close, being disentangled contributes only a small portion to the downstream performance when the downstream training steps are limited (In our case, less than or equal to 2000 steps, see Figure 2 and Figure 5).

Summary: The informativeness metric is more reliable in predicting downstream performance. Disentanglement provides only marginal additional advantages at the outset of downstream training.

5 Conclusion

In this paper, we demonstrate that dimension-wise disentanglement is not necessary for abstract visual reasoning. We identify that informativeness is of the most significance for downstream performance. Informativeness explains the previously claimed benefits of disentanglement. As abstract reasoning is a fundamental and indicative task, our study could have significant implications for a range of tasks.

Acknowledgments

This work is supported by the Ministry of Science and Technology of the People’s Republic of China, the 2030 Innovation Megaprojects “Program on New Generation Artificial Intelligence” (Grant No. 2021AAA0150000). This work is also supported by the National Key R&D Program of China (2022ZD0161700).

References

- Barrett, D.; Hill, F.; Santoro, A.; Morcos, A.; and Lillicrap, T. 2018. Measuring abstract reasoning in neural networks. In *International conference on machine learning*, 511–520. PMLR.
- Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8): 1798–1828.
- Bengio, Y.; LeCun, Y.; et al. 2007. Scaling learning algorithms towards AI. *Large-scale kernel machines*, 34(5): 1–41.
- Burgess, C.; and Kim, H. 2018. 3D Shapes Dataset. <https://github.com/deepmind/3dshapes-dataset/>.
- Burgess, C. P.; Higgins, I.; Pal, A.; Matthey, L.; Watters, N.; Desjardins, G.; and Lerchner, A. 2018. Understanding disentangling in beta-VAE. *arXiv preprint arXiv:1804.03599*.
- Cao, J.; Nai, R.; Yang, Q.; Huang, J.; and Gao, Y. 2022. An Empirical Study on Disentanglement of Negative-free Contrastive Learning. *arXiv preprint arXiv:2206.04756*.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging Properties in Self-Supervised Vision Transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*.
- Chen, R. T. Q.; Li, X.; Grosse, R.; and Duvenaud, D. 2018. Isolating Sources of Disentanglement in Variational Autoencoders. In *Advances in Neural Information Processing Systems*.
- Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, 2180–2188.
- Chen, X.; and He, K. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15750–15758.
- Comon, P. 1994. Independent component analysis, a new concept? *Signal processing*, 36(3): 287–314.
- Dittadi, A.; Träuble, F.; Locatello, F.; Wuthrich, M.; Agrawal, V.; Winther, O.; Bauer, S.; and Schölkopf, B. 2020. On the Transfer of Disentangled Representations in Realistic Settings. In *International Conference on Learning Representations*.
- Do, K.; and Tran, T. 2019. Theory and Evaluation Metrics for Learning Disentangled Representations. In *International Conference on Learning Representations*.
- Eastwood, C.; Nicolicioiu, A. L.; von Kügelgen, J.; Kekić, A.; Träuble, F.; Dittadi, A.; and Schölkopf, B. 2022. DCI-ES: An Extended Disentanglement Framework with Connections to Identifiability. *arXiv preprint arXiv:2210.00364*.
- Eastwood, C.; and Williams, C. K. 2018. A framework for the quantitative evaluation of disentangled representations. In *ICML*.
- Fleuret, F.; Li, T.; Dubout, C.; Wampler, E. K.; Yantis, S.; and Geman, D. 2011. Comparing machines and humans on a visual categorization test. *Proceedings of the National Academy of Sciences*, 108(43): 17621–17625.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep learning*. MIT press.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P. H.; Buchatskaya, E.; Doersch, C.; Pires, B. A.; Guo, Z. D.; Azar, M. G.; et al. 2020. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*.
- Hahne, L.; Lüddecke, T.; Wörgötter, F.; and Kappel, D. 2019. Attention on abstract visual reasoning. *arXiv preprint arXiv:1911.05990*.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16000–16009.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- Higgins, I.; Amos, D.; Pfau, D.; Racaniere, S.; Matthey, L.; Rezende, D.; and Lerchner, A. 2018. Towards a definition of disentangled representations. *arXiv preprint arXiv:1812.02230*.
- Higgins, I.; Chang, L.; Langston, V.; Hassabis, D.; Summerfield, C.; Tsao, D.; and Botvinick, M. 2021. Unsupervised deep learning identifies semantic disentanglement in single inferotemporal face patch neurons. *Nature communications*, 12(1): 6456.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2016. beta-vae: Learning basic visual concepts with a constrained variational framework. *arXiv preprint*.
- Jeon, I.; Lee, W.; and Kim, G. 2018. IB-GAN: Disentangled representation learning with information bottleneck GAN. *arXiv preprint*.
- Kim, H.; and Mnih, A. 2018. Disentangling by factorising. In *International Conference on Machine Learning*, 2649–2658. PMLR.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.

- Kumar, A.; Sattigeri, P.; and Balakrishnan, A. 2017. Variational inference of disentangled latent concepts from unlabeled observations. *arXiv preprint arXiv:1711.00848*.
- Lake, B. M.; Ullman, T. D.; Tenenbaum, J. B.; and Gershman, S. J. 2017. Building machines that learn and think like people. *Behavioral and brain sciences*, 40.
- Leeb, F.; Lanzillotta, G.; Annadani, Y.; Besserve, M.; Bauer, S.; and Schölkopf, B. 2020. Structure by architecture: disentangled representations without regularization. *arXiv preprint arXiv:2006.07796*.
- Lin, Z.; Thekumparampil, K.; Fanti, G.; and Oh, S. 2020. Infogan-cr and modelcentrality: Self-supervised model training and selection for disentangling gans. In *International Conference on Machine Learning*, 6127–6139. PMLR.
- Locatello, F.; Abbati, G.; Rainforth, T.; Bauer, S.; Schölkopf, B.; and Bachem, O. 2019a. On the fairness of disentangled representations. *Advances in Neural Information Processing Systems*, 32.
- Locatello, F.; Bauer, S.; Lucic, M.; Raetsch, G.; Gelly, S.; Schölkopf, B.; and Bachem, O. 2019b. Challenging common assumptions in the unsupervised learning of disentangled representations. In *ICML*.
- Locatello, F.; Poole, B.; Rätsch, G.; Schölkopf, B.; Bachem, O.; and Tschannen, M. 2020. Weakly-supervised disentanglement without compromises. In *International Conference on Machine Learning*, 6348–6359. PMLR.
- Małkiński, M.; and Mańdziuk, J. 2022. Deep Learning Methods for Abstract Visual Reasoning: A Survey on Raven’s Progressive Matrices. *arXiv preprint arXiv:2201.12382*.
- Matthey, L.; Higgins, I.; Hassabis, D.; and Lerchner, A. 2017. dSprites: Disentanglement testing Sprites dataset. <https://github.com/deepmind/dsprites-dataset/>.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. DINOv2: Learning Robust Visual Features without Supervision. *arXiv preprint arXiv:2304.07193*.
- Peters, J.; Janzing, D.; and Schölkopf, B. 2017. *Elements of causal inference: foundations and learning algorithms*. The MIT Press.
- Raven, J. C. 1941. Standardization of progressive matrices, 1938. *British Journal of Medical Psychology*.
- Ridgeway, K.; and Mozer, M. C. 2018. Learning deep disentangled embeddings with the f-statistic loss. *arXiv preprint arXiv:1802.05312*.
- Roth, K.; Ibrahim, M.; Akata, Z.; Vincent, P.; and Bouchacourt, D. 2022. Disentanglement of Correlated Factors via Hausdorff Factorized Support. *arXiv preprint arXiv:2210.07347*.
- Schmidhuber, J. 1992. Learning factorial codes by predictability minimization. *Neural computation*, 4(6): 863–879.
- Schölkopf, B.; Locatello, F.; Bauer, S.; Ke, N. R.; Kalchbrenner, N.; Goyal, A.; and Bengio, Y. 2021. Toward causal representation learning. *Proceedings of the IEEE*, 109(5): 612–634.
- Steenbrugge, X.; Leroux, S.; Verbelen, T.; and Dhoedt, B. 2018. Improving generalization for abstract reasoning tasks using disentangled feature representations. *arXiv preprint arXiv:1811.04784*.
- Träuble, F.; Dittadi, A.; Wuthrich, M.; Widmaier, F.; Gehler, P. V.; Winther, O.; Locatello, F.; Bachem, O.; Schölkopf, B.; and Bauer, S. 2021. The Role of Pretrained Representations for the OOD Generalization of RL Agents. In *International Conference on Learning Representations*.
- Tschannen, M.; Bachem, O.; and Lucic, M. 2018. Recent advances in autoencoder-based representation learning. *arXiv preprint arXiv:1812.05069*.
- van Steenkiste, S.; Locatello, F.; Schmidhuber, J.; and Bachem, O. 2019. Are disentangled representations helpful for abstract visual reasoning? *arXiv:1905.12506*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, Y.; and Jordan, M. I. 2021. Desiderata for representation learning: A causal perspective. *arXiv preprint arXiv:2109.03795*.
- Whittington, J. C.; Dorrell, W.; Ganguli, S.; and Behrens, T. E. 2022. Disentangling with Biological Constraints: A Theory of Functional Cell Types. *arXiv preprint arXiv:2210.01768*.