

# Learning Interpretable Concepts: Unifying Causal Representation Learning and Foundation Models

Goutham Rajendran<sup>\*1</sup>, Simon Buchholz<sup>\*2,3</sup>,  
 Bryon Aragam<sup>4</sup>, Bernhard Schölkopf<sup>2,5</sup>, and Pradeep Ravikumar<sup>1</sup>

<sup>1</sup>Machine Learning Dept., Carnegie Mellon University, Pittsburgh, USA

<sup>2</sup>Max Planck Institute for Intelligent Systems, Tübingen, Germany

<sup>3</sup>Tübingen AI Center

<sup>4</sup>University of Chicago, Chicago, USA

<sup>5</sup>ELLIS Institute Tübingen

December 10, 2024

## Abstract

To build intelligent machine learning systems, there are two broad approaches. One approach is to build inherently interpretable models, as endeavored by the growing field of causal representation learning. The other approach is to build highly-performant foundation models and then invest efforts into understanding how they work. In this work, we relate these two approaches and study how to learn human-interpretable concepts from data. Weaving together ideas from both fields, we formally define a notion of concepts and show that they can be provably recovered from diverse data. Experiments on synthetic data and large language models show the utility of our unified approach.

## 1 Introduction

A key goal of modern machine learning is to learn representations of complex data that are human-interpretable and can be controlled. This goal is of paramount importance given the breadth and importance of ML in today’s world. There seem to be two broad approaches toward such intelligent systems. The first approach is to build models that are inherently interpretable and then subsequently focus on how to extract maximum performance from them; and the second approach is to build high-performance neural models, and then subsequently invest efforts to understand the inner workings of such models.

A prominent example of the first camp is the field of Causal Representation Learning (CRL) [82, 81]. CRL is an intricate interplay of ideas from causality, latent variable modeling and deep learning, with the main goal being to reconstruct the true generative factors of data. To ensure that the true generative factors can be provably recovered, CRL relies on the central theme of *identifiability* which posits that a unique model fits the data, which in turn implies that the problem of learning the generative factors is well-posed and therefore should theoretically be amenable to modern techniques. If such a generative model reconstruction can be done, the model will naturally enjoy a host of desired properties such as robustness and generalization. While this endeavor has been somewhat successful in many domains such as computer vision [41, 103, 2], robotics [58, 9, 54, 113] and genomics [90, 112], it is unclear how it relates to the research on foundation models.

The other camp is more empirical, where one tries to build a high-performance model where performance is measured via various downstream tasks and then eventually invest efforts into explaining or interpreting how they work. For instance, large language models and other foundation models are built to be highly performant for a variety of tasks. Owing to their incredible success, there is a growing but heavily-debated belief that such models are truly “intelligent” because they have indeed learned the true underlying generative factors somehow, sometimes referred to as the “world model”. While we are far from scientifically verifying this, the community has invested tremendous efforts into interpretability research of foundation

---

<sup>\*</sup>Equal Contribution

models, e.g., the field of mechanistic interpretability [66] aims to reverse engineer what large language models learn.

In this work, we take a first step toward unifying these approaches. We focus on the goal of learning identifiable human-interpretable concepts from complex high-dimensional data. Specifically, we build a theory of what concepts mean for complex high-dimensional data and then study under what conditions such concepts are identifiable, i.e., when can they be unambiguously recovered from data. To formally define concepts, we leverage extensive empirical evidence in the foundation model literature that surprisingly shows that, across multiple domains, human-interpretable concepts are often *linearly* encoded in the latent space of such models (see Section 2), e.g., the sentiment of a sentence is linearly represented in the activation space of large language models [96]. Motivated by this rich empirical literature, we formally define concepts as affine subspaces of some underlying representation space. Then we connect it to causal representation learning by proving strong identifiability theorems for *only desired concepts* rather than all possible concepts present in the true generative model. Therefore, in this work we tread the fine line between the rigorous principles of causal representation learning and the empirical capabilities of foundation models, effectively showing how causal representation learning ideas can be applied to foundation models.

Let us be more concrete. For observed data  $X$  that has an underlying representation  $Z_u$  with  $X = f_u(Z_u)$  for an arbitrary distribution on  $Z_u$  and a (potentially complicated) nonlinear underlying mixing map  $f_u$ , we define concepts as affine subspaces  $AZ_u = b$  of the latent space of  $Z_u$ s, i.e., all observations falling under a concept satisfy an equation of this form. Since concepts are not precise and can be fuzzy or continuous, we will allow for some noise in this formulation by working with the notion of concept conditional distributions (Definition 3). Of course, in general,  $f_u$  and  $Z_u$  are very high-dimensional and complex, as they can be used to represent arbitrary concepts. Instead of ambitiously attempting to reconstruct  $f_u$  and  $Z_u$  as CRL would do, we go for a more relaxed notion where we attempt to learn a minimal representation that represents *only the subset of concepts we care about*; i.e., a simpler decoder  $f$  and representation  $Z$ —different from  $f_u$  and  $Z_u$ —such that  $Z$  linearly captures a subset of relevant concepts as well as a valid representation  $X = f(Z)$ . With this novel formulation, we formally prove that concept learning is identifiable up to simple linear transformations (the linear transformation ambiguity is unavoidable and ubiquitous in CRL). This relaxes the goals of CRL to only learn relevant representations and not necessarily learn the full underlying model. It further suggests that foundation models do in essence learn such relaxed representations, partially explaining their superior performance for various downstream tasks.

Apart from the above conceptual contribution, we also show that to learn  $n$  (atomic) concepts, we only require  $n + 2$  environments under mild assumptions. Contrast this with the adage in CRL [41, 11] where we require  $\dim(Z_u)$  environments for most identifiability guarantees, where as described above we typically have  $\dim(Z_u) \gg n + 2$ . These theoretical insights are then validated on synthetic data, where we use a contrastive algorithm to learn such representations for a given collection of concepts.

Moving ahead to real-world data and foundation models, we proceed to show an effective application of our framework to large language models (LLMs). In particular, we consider the alignment problem of making pre-trained LLMs more *truthful*. First, we make the assumption that pre-trained LLMs have already learnt the concept of truth linearly, as has been empirically observed in Li et al. [51] (see also Section 2 for more evidence). Therefore, we can apply our ideas, which we use to mechanistically reason about how the recent Inference-Time Intervention technique [51] steers pre-trained LLMs to give more truthful responses to questions. Then, we exploit our insights to extend this technique, by building steering matrices rather than steering vectors to align LLMs towards concepts. Preliminary experiments using LLaMA [97] show the efficacy of our approach on the TruthfulQA dataset [53].

In summary, our contributions are:

1. We formalize the notion of distributions induced by abstract concepts in complex domains such as images or text. Our definition of concept conditional distributions allows both continuous and fuzzy concepts.
2. We prove near-optimal identifiability results for learning a collection of concepts from a diverse set of environments. We also verify our guarantees via a contrastive learning algorithm on synthetic data. Thus, our work presents a novel framework for identifying concepts by weaving together ideas from causal representation learning and foundation models.
3. We show the applicability of our ideas toward mechanistic interpretability, by explaining why inference-time steering vectors align large language models toward abstract concepts such as

truthfulness. Furthermore, we extend this idea and propose to use steering matrices instead of steering vectors for better alignment. Our experiments with LLaMA [97] on TruthfulQA [53] show improved performance.

## 2 Related work

**Causal representation learning** Causal representation learning (CRL) [82, 81] aims to learn generative factors of high-dimensional data. This exciting field has seen significant progress in the last few years [41, 9, 85, 47, 62, 45, 11, 28, 1, 104]. A fundamental perspective in this field is to ensure that the model parameters we attempt to recover are identifiable [41, 18, 106]. Identifiability is the notion that the model parameters we learn are equivalent to the true model parameters up to simple transformations. However, it's not clear until this work how this relates to the representations learned by foundation models. Our work acts as a bridge between these two approaches. We will elaborate more on the connection of our framework to CRL in Appendix B.

**Linearity of representations** Sometimes referred to as the linear representation hypothesis, it is commonly believed that well-trained foundation models in multiple domains learn linear representations of human-interpretable concepts, with experimental evidence going back at least a decade [61, 92, 4]. This has been experimentally observed in computer vision models [72, 76, 7, 23, 43, 107, 98], language models [61, 70, 4, 16, 95, 22], large language models [12, 96, 65, 63, 51, 68, 30, 40], and other intelligent systems [59, 83]. Various works have also attempted to justify why this happens [49, 4, 27, 3, 24, 84].

We take a different angle: Given that this phenomenon has been observed for certain concepts of interest, how does this enable recovery of the concepts themselves? Consequently, our model assumptions are well-founded and our theory applies to multiple domains of wide interest.

**Concepts from pre-trained models** Concept discovery is an important sub-field of machine learning which attempts to understand the behavior of powerful neural models. We do not attempt to list the numerous experimental works in this direction, see e.g., Schut et al. [83], Carvalho et al. [13]. However, theoretical progress in this direction is relatively limited. Prior works have attempted to formalize the notion of concepts [107, 68, 83], however their definitions seem specific to the model and domain under consideration, e.g., Park et al. [68] focus on binary concepts via large language model representations of counterfactual word pairs, whereas our general concept definitions are applicable to all domains. Importantly, they do not provide identifiable recovery guarantees from data. To the best of our knowledge, ours is the first work to do this for general human-interpretable concepts, by exploiting ideas from the causal representation learning literature.

## 3 Setup

In this section, we provide a formal definition of concepts, which are high-level abstractions present in data. This allows us to develop a theoretical framework for associated data distributions and identifiability theory. For the sake of intuition, we can think of the data as images of different objects and the color of the object as a concept.

### 3.1 Generative model

We assume that the observed data  $X$  lies in a space  $\mathcal{X} \subseteq \mathbb{R}^{d_x}$  of dimension  $d_x$  and has an underlying representation  $X = f(Z)$  for latent variables  $Z$  that lie in a latent concept space  $\mathbb{R}^{d_z}$  of dimension  $d_z$ . In contrast to most prior works we do not take the viewpoint that  $Z$  represents the true underlying mechanism that generated the data. Instead we simply assume that the latent representation has the geometric property that it maps certain regions of the observation space to linear subspaces of the latent space (motivated by the foundation models literature). Then we investigate to what extent this assumption results in the identifiability of such representations.

**Assumption 1** (Mixing function). *The non-linear  $f$  is injective and differentiable.*

Injectivity and differentiability are standard assumptions in representation learning. However, we make no additional assumptions on  $f$ : The map from latent space to observation space can be arbitrarily non-linear. For the rest of this subsection, fix a data representation  $f^{-1}$ .

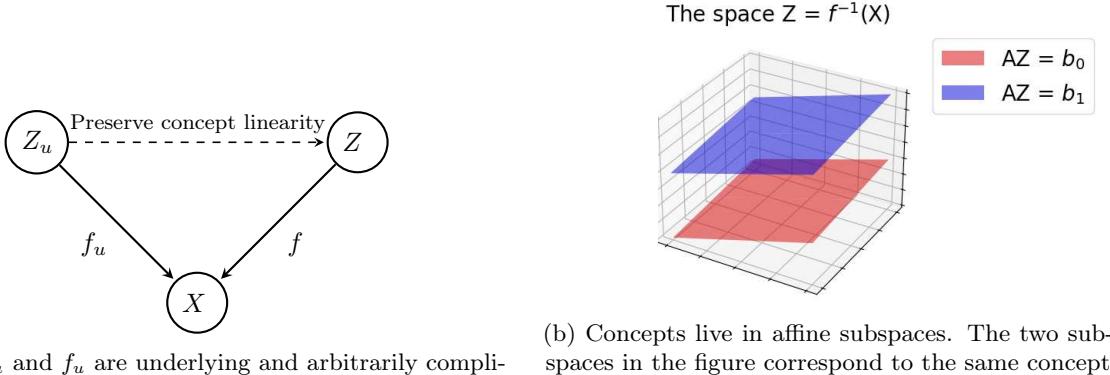


Figure 1: Illustration of our framework

We now define concepts living in the latent space  $\mathbb{R}^{d_z}$ . Informally, we think of an (atomic) concept as a vector  $a \in \mathbb{R}^{d_z}$  such that for a given valuation  $b \in \mathbb{R}^n$ , the set of all observations  $X$  that satisfy this concept is given by  $\{X = f(Z) | \langle a, Z \rangle = b\}$ . For instance, for an object in an image  $X$ , if  $a \in \mathbb{R}^{d_z}$  is the concept of red color,  $b \in \mathbb{R}$  could indicate the intensity; then all datapoints  $X$  satisfying this concept, i.e., all images with an object that has color red with intensity  $b$ , can be characterized as  $X = f(Z)$  where  $Z$  satisfies  $\langle a, Z \rangle = b$ . For a 3D visualization, see Fig. 1b. We make this intuition formal below.

**Definition 1** (Concepts). *A concept  $C$  is a linear transformation  $A : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_C}$ . The dimension of the concept will be denoted by  $\dim(C) = d_C$ . A valuation is a vector  $b \in \mathbb{R}^{d_C}$  and we say that a datapoint  $X$  satisfies the concept  $C$  with valuation  $b$  if the evaluation map satisfies  $AZ = b$  where  $Z = f^{-1}(X)$ .*

In this work, we are interested in learning a collection of  $m$  concepts  $C^1, \dots, C^m$  from observed data. By left multiplying by the pseudo-inverse  $A^+$ , we can equivalently assume  $A$  is a projector matrix. However, the current definition is more suitable for embeddings of real models.

When we talk of learning concepts  $C$ , we are in particular interested in learning the evaluation map  $Af^{-1}(x)$ . This is a more modest objective than learning the entire map  $f$  which is the usual goal in identifiability theory. While the latter typically requires stringent assumptions, in particular  $\Omega(d_z)$  environments are necessary, our weaker identifiability results only need  $O(d_C) \ll O(d_z)$  environments. To simplify our analysis, we can view each concept as being composed of one dimensional concepts, which warrants the following definition.

**Definition 2** (Atoms). *We define an atom (short for atomic concept) to be any concept  $C$  with  $\dim(C) = 1$ .*

Intuitively, atomic concepts are fundamental concepts that live in a space of co-dimension 1 in latent space, and thus are equivalently defined by vectors  $a \in \mathbb{R}^{d_z}$ . For the sake of intuition, we can think of red color, size of object, etc., as examples of atomic concepts. Any generic concept is then composed of a collection of atomic concepts, e.g., the concept  $C$  of all small dark red objects will correspond to  $\dim(C) = 2$  with row 1 corresponding to the atomic concept of red color with large valuation (dark red objects) and row 2 corresponding to the atomic concept of object size with low valuation (small objects).

### 3.2 Data distributions

We now define the distributions of datasets over concepts. We will predominantly work with distributions of  $Z$  over  $\mathbb{R}^{d_z}$ , as the resulting distribution of  $X = f(Z)$  over  $\mathbb{R}^{d_x}$  can be obtained via a simple change of variables.

To build intuition, consider the case where we first collect a base dataset with some underlying distribution and then collect concept datasets via filtering. For instance, we could first collect a set of images of all objects and then, to collect a dataset of dark red colored objects, we filter them to only keep images of dark red colored objects. We call the former the *base distribution* and the latter the *concept conditional distribution* corresponding to our concept.

Fix a nonlinearity  $f$ . We assume that the base data distribution is the distribution of  $X = f(Z)$  with  $Z \sim p$ , where  $p$  is the underlying distribution on  $\mathbb{R}^{d_z}$ . In what follows, we will abuse notation and use  $p$  for both the distribution and the corresponding probability density which we assume exists. We make no further assumptions on  $p$  since we do not wish to model the collection of real-life datasets that have been collected from nature and which could be very arbitrary.

We now define the concept conditional distribution, which is a distribution over  $X$  that is induced by noisy observations of a particular concept at a particular valuation. Formally, assume we want to condition on some atomic concept  $a \in \mathbb{R}^{d_z}$  with valuation  $b$ . It is reasonable to assume that this conditioning is a noisy operation. For instance, human beings are great at distilling concepts from noised images, e.g., they recognize cars in a misty environment. We formalize this by assuming that data collection is based on a noisy estimate  $\tilde{b} = \langle a, z \rangle + \epsilon$  where  $\epsilon$  is independent of  $z$  and its density is a symmetric distribution with density  $q(\epsilon)$ . Then we consider the concept conditional distribution

$$\begin{aligned} p_C(z) &= p(z|\tilde{b} = b) \propto p(\tilde{b} = b|z)p(z) \\ &= q(b - \langle a, z \rangle)p(z) \end{aligned} \tag{1}$$

where we used Bayes theorem in the last step. This definition directly extends to higher dimensional concepts which are concisely defined as follows.

**Definition 3** (Concept conditional distribution). *For a concept  $C$  with associated linear map  $A$  and an arbitrary valuation  $b \in \mathbb{R}^{\dim(C)}$ , we define the concept conditional distribution to be the set of observations  $X$  respecting this concept, which is defined as the resampled distribution of  $X = f(Z)$  where  $Z \sim p_C$  with*

$$p_C(Z) \propto p(Z) \prod_{k \leq \dim(C)} q((AZ - b)_k) \tag{2}$$

This is by no means the only possible definition, and we present feasible alternate definitions in Appendix C. We remark that our formulation is related to the iVAE setting [41] and the auxiliary variable setting for identifiable ICA in Hyvarinen et al. [36] and we discuss the relation later. Note that the majority of recent identifiability results relied on interventional data while we only consider conditional information here.

We are ready to define our main problem of interest.

**Problem 1.** *We are given an observational dataset  $X^0 = f(Z^0)$  corresponding to the latent base distribution  $p$  along with datasets  $X^1, \dots, X^m$  corresponding to concept conditional datasets for different concepts  $C^1, \dots, C^m$  and corresponding valuations  $b^1, \dots, b^m$  over the same latent space  $\mathbb{R}^{d_z}$  with the same mixing  $f$ . Under what conditions (and up to which symmetries) can we learn the concepts  $C^1, \dots, C^m$ , which includes the linear maps  $A^1, \dots, A^m$ , and the concept valuations  $A^e f^{-1}(x)$ ?*

Toward this end, a fundamental question is whether this problem is even possible, i.e., whether it is well-defined. This ties to the question of identifiability of such parameters of interest. Therefore, we make the following definition. Informally, for the setting above, we say that the concepts  $(C^1, A^1), \dots, (C^m, A^m)$  with associated nonlinearity  $f$  are identifiable (and thus learnable) if for any other collection of different parameters that fit the data, they are linearly related to the true parameters.

**Definition 4** (Identifiability). *Given datasets  $X^0, X^1, \dots, X^m$  corresponding to the observational distribution and  $m$  concepts  $C^1, \dots, C^m$  with underlying distribution  $p$  on  $\mathbb{R}^{d_z}$ , nonlinearity  $f$ , linear maps  $A^1, \dots, A^m$  and valuations  $b^1, \dots, b^m$ , we say the concepts are identifiable if the following holds: Consider any different collection of parameters  $\tilde{f}, \tilde{d}_z, \tilde{p}$ , concepts  $(\tilde{C}^1, \tilde{A}^1), \dots, (\tilde{C}^m, \tilde{A}^m)$  and valuations  $\tilde{b}^1, \dots, \tilde{b}^m$  that also generate the same observations  $X^0, X^1, \dots, X^m$ . Then there exists a shift  $w \in \mathbb{R}^{d_z}$ , permutation matrices  $P^e$  and invertible diagonal matrices  $\Lambda^e$  such that for all  $e$  and  $x$ ,*

$$\tilde{A}^e \tilde{f}^{-1}(x) = \Lambda^e P^e A^e (f^{-1}(x) + w), \tag{3}$$

i.e., we can evaluate the concept evaluations on the data up to linear reparametrizations. Moreover, there exists a linear map  $T : \mathbb{R}^{\tilde{d}_z} \rightarrow \mathbb{R}^{d_z}$  such that the concepts and their evaluations are related as follows

$$\tilde{A}^e = P^e A^e T^{-1}, \quad \tilde{b}^e = \Lambda^e P^e (b^e - A^e w). \tag{4}$$

The main message of this definition is that identifiability implies we can identify the nonlinear map  $f^{-1}$  within the span of the subspace of the concepts of interest, and therefore we can recover the concepts of interest from our data. That is, if certain concepts are identifiable, then we will be able to learn these concept representations up to linearity, even if they can be highly nonlinear functions of our data. Such concept discovery is useful because they can then be used for further downstream tasks such as controllable generative modeling.

We emphasize that in contrast to previous work we are not aiming to identify  $f$  completely. Let us now clarify why no stronger identifiability results can be expected. Firstly, we cannot hope to resolve the linear transformation ambiguity because the latent space is not directly observed. In other words, a concept evaluation can be defined either as  $\langle a, Z \rangle$  or as  $\langle Ta, T^{-\top}Z \rangle$  for an invertible linear map  $T$ . However, for the purposes of downstream tasks, this is fine since the learned concepts will still be the same in principle. Secondly, we cannot expect to recover  $f^{-1}$  outside the span of the concepts because we do not manipulate the linear spaces outside the span therefore we do not learn this information from our observed data so this is also tight. The permutation matrix captures the fact that the ordering of the concepts does not matter. Therefore, this definition captures the most general identifiability guarantee that we can hope for in our setting and furthermore, this suffices for downstream tasks such as controllable data generation.

Because we will only be interested in recovering the set of concepts up to linear transformations, without loss of generality, we will fix the base collection of atomic concepts. That is, we assume that each concept  $C^e$  corresponds to a linear map  $A^e$  whose rows are a subset of  $\mathcal{C}$ , where  $\mathcal{C} = \{a_1, \dots, a_n\}$  is a set of atomic concepts that we wish to learn. Moreover, we assume that they are linearly independent, since we want them to encode distinct concepts.

**Assumption 2.** *There exists a set of atomic concepts  $\mathcal{C} = \{a_1, \dots, a_n\}$  of linearly independent vectors such that each concept  $C^e$  under consideration contains as rows a subset  $S^e$  of  $\mathcal{C}$  and all concepts in  $\mathcal{C}$  appear, i.e.,  $\bigcup_e S^e = [n]$ .*

**Remark 1.** *Note that identifiability as defined in Definition 4 implies that the atoms can be identified in the sense that there is a permutation  $\pi \in S_n$  and  $\lambda_i \neq 0$  such that for  $T$  as in Definition 4 and some  $\lambda_i$*

$$\tilde{a}_{\pi(i)}^\top = a_i^\top T^{-1} \quad (5)$$

$$\langle \tilde{a}_{\pi(i)}, \tilde{f}^{-1}(x) \rangle = \lambda_i (\langle a_i, f^{-1}(x) \rangle + \langle a_i, w \rangle), \quad (6)$$

i.e., we can evaluate the valuations of the atomic concepts up to linear reparametrization.

## 4 Main Result

In this section, we will be interested to identify a class of distinct concepts from data. As discussed earlier, we wish to recover them as per our precise definition of identifiability, which as we saw is the best possible.

The punchline is that when we have rich datasets, i.e., sufficiently rich concept conditional datasets, then we can recover the concepts. Importantly, we only require a number of datasets that depends only on the number of atoms  $n$  we wish to learn (in fact,  $O(n)$  datasets), and not on the underlying latent dimension  $d_z$  of the true generative process. This is a significant departure from most works on causal representation learning, since the true underlying generative process could have  $d_z = 1000$ , say, whereas we may be interested to learn only  $n = 5$  concepts, say. In this case, causal representation learning necessitates at least  $\sim 1000$  datasets, whereas we show that  $\sim n + 2 = 7$  datasets are enough if we only want to learn the  $n$  atomic concepts. We will explain the connection to causal representation learning in Appendix B. Let us now discuss our main assumptions.

**Assumption 3.** *We choose the distribution  $q$  to be Gaussian with mean 0 and variance  $\sigma^2$  for some  $\sigma^2 > 0$ .*

We relate the concepts  $C^e$  to the atoms. Recall that we defined the index sets  $S^e = \{i \in [n] : a_i \in \mathcal{C}$  is a row of  $A^e\}$  of atomic concepts in environment  $e$ .

We define the environment-concept matrix  $M \in \mathbb{R}^{m \times n}$  indexed by environments and atoms by

$$M_{ei} = \begin{cases} \frac{1}{\sigma^2} & \text{if } i \in S^e \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

Similarly, we consider the environment-valuation matrix  $B \in \mathbb{R}^{m \times n}$  given by

$$B_{ei} = \begin{cases} \frac{b_k^e}{\sigma^2} & \text{if } i \in S^e \text{ and row } k \text{ of } A^e \text{ is } a_i, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

We now state our assumptions on the diversity of the concept conditional distributions that will ensure identifiability.

**Assumption 4** (Environment diversity I). *The environment-concept matrix  $M \in \mathbb{R}^{m \times n}$  has rank  $n$  and there is a vector  $v \in \mathbb{R}^m$  such that  $v^\top M = 0$  and all entries of  $v^\top B$  are non-zero ( $B$  denotes that environment-valuation matrix).*

We remark that this assumption can only hold for  $m \geq n + 1$  and indeed is satisfied under mild assumptions on the environments if  $m = n + 1$ , as the following lemma shows.

**Lemma 1.** *Assumption 4 is satisfied almost-surely if there are  $n + 1$  concept conditional distributions such that every  $n$  rows of the environment-concept matrix are linearly independent and the  $b^e$  are drawn independently according to a continuous distribution.*

We also assume one additional diversity condition.

**Assumption 5** (Environment diversity II). *For every pair of atoms  $a_i$  and  $a_j$  with  $i \neq j$  there is an environment  $e$  such that  $i \in S^e$  and  $j \notin S^e$ .*

We remark that these are the only assumptions about the sets  $S^e$ . In particular, we do not need to know the sets  $S^e$ . In the proof, we will extract these sets based on the signatures they leave on the datasets. We can now state our main result.

**Theorem 1.** *Suppose we are given  $m$  context conditional datasets  $X^1, \dots, X^m$  and the observational dataset  $X^0$  such that Assumptions 1, 2, 3, 4, and 5 hold. Then the concepts are identifiable as per Definition 4.*

**Remark 2.** *Assumption 4 can only be satisfied for  $m \geq n + 1$ , i.e., the result requires at least  $n + 2$  environments. On the other hand, Lemma 1 assures that  $n + 2$  environments are typically sufficient. We expect that the result could be slightly improved by showing identifiability for  $n + 1$  environments under suitable assumptions. However, this would probably require more advanced techniques from algebraic statistics [20] compared to the necessary techniques for our result.*

As mentioned before, our setting somewhat resembles the iVAE setting in Khemakhem et al. [41] and therefore, their proof techniques can also be applied, with several modifications, to derive identifiability results in our setting (however our formulation and application are very different). We notably remark that this approach will require more environments because their main assumption is that the matrix  $\Lambda = (M, B) \in \mathbb{R}^{m \times 2n}$  has rank  $2n$  so that  $2n + 1$  environments are necessary. Moreover, this rank condition is much stronger than Assumption 4. For completeness and as a warm-up we prove this result in Appendix A.

The full proof of Theorem 1 can be found in Appendix A. Although the proof is fairly involved, we present the core idea here. Given two sets of parameters that fit the data, we show they're linearly related. To do this, we first equate the log-odds, which by assumption will be quadratic polynomials. Taking derivatives, we will be able to obtain linear equations with coefficients depending on our environment matrices. By a careful linear algebraic analysis, we can conclude identifiability. The diversity assumptions ensure invertibility of various matrices that show up, which lets us solve the equations.

## 5 Experiments

We present two slates of experiments to supplement our framework. In Section 5.1, we validate our results on synthetic data, via an end-to-end contrastive learning algorithm for concept learning. We additionally propose a rejection sampling algorithm to sample from concept conditional distributions in Appendix G. In Section 5.2, we focus on real-world settings especially on large language models (LLMs). We show how our techniques can be used to align pre-trained LLMs towards abstract concepts such as truthfulness.

### 5.1 End-to-end Contrastive learning algorithm

In this section, we will validate our insights on synthetically generated data in order to verify our results on identifiability. We present an end-to-end framework based on contrastive learning (similar to [11]) to learn the nonlinearity as well as concepts from data. The model architecture is designed based on our concept conditional distribution parametrization.

The core idea is to utilize contrastive learning as follows. For each concept conditional distribution  $X^e$ , we train a neural network to distinguish concept samples  $x \sim X^e$  from base samples  $x \sim X^0$ . First, we compute the true log-odds for this classification problem.

**Lemma 2.** For any concept index  $e$ , there exist some constants  $c_e$  such that

$$\begin{aligned} & \ln(p^e(Z)) - \ln(p(Z)) \\ &= \sum_{i=1}^n \left( -\frac{1}{2} M_{ei} \langle a_i, Z^e \rangle^2 + B_{ei} \langle a_i, Z^e \rangle \right) + c_e \end{aligned}$$

where  $M, B$  are the environment-concept matrix and the environment-valuation matrix defined in (7) and (8).

The proof is in Appendix E. Now, let's try to learn the  $n$  atomic concepts up to linearity. To do this, we build a neural architecture for this classification problem with the final layer mimicking the log-odds expression above, which can then be trained end-to-end. Because of the careful parametrization of the last layer, this will encourage the model to learn the representations as guaranteed by our results. The details are deferred to Appendix E.

**Sampling from concept conditional distributions** A common task in controllable generative modeling is being able to generate data from a known concept. Note that this is not straightforward in our setting because the normalization term in Eq. (2) is not efficiently computable. To do this efficiently, we also outline a simple algorithm (Algorithm 1 in Appendix G) to sample from the concept conditional distribution for a known concept. Our proposed algorithm is based on rejection sampling and the algorithm as well as the complexity analysis are deferred to Appendix G.

**Synthetic experiments** We test our proposed method on synthetic data as follows. We sample the base distribution from a Gaussian Mixture model and experiment with both linear and nonlinear mixing functions. The parameters are all chosen randomly. The number of concepts  $n$  is intentionally chosen to be less than the ground truth dimension  $d_z$  and the number of concepts is  $m = n + 1$  as per our theory. Additional details are deferred to Appendix F.

Table 1: Linear identifiability when number of concepts  $n$  is less than underlying latent dimension  $d_z$ , averaged over 5 seeds

Type of mixing $f$	$n$	$d_z$	Obs dim $d_x$	$R^2 \uparrow$	MCC $\uparrow$
Linear	2	3	4	0.98	0.98
Nonlinear	2	3	4	0.94	0.96
Linear	3	4	6	0.99	0.86
Nonlinear	3	4	6	0.97	0.92
Linear	4	8	10	0.97	0.87
Nonlinear	4	8	10	0.94	0.87

In Table 1, we show the metrics achieved by our contrastive learning algorithm on this data, where we report the  $R^2$  and Mean Correlation Coefficient (MCC) metrics [41, 42] of the recovered latents with respect to the ground truth concept valuations. There are no baselines since we are in a novel setting where standard causal representation learning baselines do not easily apply. However, our metrics are comparable to what's usually reported in such highly nonlinear settings [109, 11].

We remark that variations of the contrastive method can be designed for harder synthetic settings and different problems related to concept discovery. However, we will move onto real-life data experiments next.

## 5.2 Alignment of Large Language Models

In this section, we show an application of our framework to alignment of Large Language Models (LLMs). In particular, we exploit our ideas to improve the Inference-Time Intervention technique [51] to promote LLMs to be more truthful. In contrast to previous sections, we will focus on exactly one concept – *truthfulness* with two relevant valuations – *false* and *true*. We concretize this by assuming this concept is linear and the valuations correspond to real numbers (as evidenced by Li et al. [51], see also Section 2 for more evidence). The downstream task is to make pre-trained LLMs answer truthfully, i.e., change the valuation of this concept from *false* to *true*, without affecting any other orthogonal concepts.

Thus, in contrast to the previous section where we trained an end-to-end model to learn both the non-linearities as well as the concepts, we will now assume that the non-linearity has already been learned up to a linear transformation (by large-scale training of LLMs). This aligns with our theoretical insights because the training data for powerful LLMs are diverse, so they essentially satisfy our core assumptions (see also the related work [29] that proposes that context is environment in LLM training). Therefore, we simply focus on the downstream tasks, which in this section is LLM alignment. The difficulty, of course, is that we do not know the concept matrix or the valuations. A detailed introduction to large language models (LLMs) and the Inference-Time Intervention (ITI) technique is deferred to Appendix D. We present a simplified summary here.

Given a context  $x$ , a GPT-style LLM outputs a representation  $f(x)$ , which is then decoded into probabilities of the next token. ITI is an activation patching technique that controls the behavior of LLMs during inference time, in particular it promotes truthfulness. ITI proposes to modify the forward pass (during inference time) to  $f(x) + \alpha\sigma \cdot \eta$  where  $\eta$  is a trained steering vector,  $\sigma$  is the standard deviation of the representations in this direction and  $\alpha$  is a tunable hyperparameter.<sup>1</sup>

Here, the steering vectors are intended to steer the LLM towards truthful responses. To train them, the work [51] uses a set of counterfactual pairs  $(c_i^F, c_i^T)$  of false and true responses to the same questions (for instance from the TruthfulQA dataset [53]), compute the difference of their representations  $f(c_i^T) - f(c_i^F)$  and the normalized mean difference is used as the steering vector  $\eta$ .

First, we use our framework to mechanistically reason about how ITI modifies the internal representations and illustrate why the mean is a good choice. We present an intuitive summary here, postponing details to Appendix D. If we consider the concept conditional distributions of the true and false valuations of our truthfulness concept (assume for simplicity it's an atomic concept  $a$  with valuations  $b^T, b^F$ ), then they concentrate in their respective hyperplanes  $\langle a, f(c_i^T) \rangle = b^T$  and  $\langle a, f(c_i^F) \rangle = b^F$ . Now, consider the steering vector  $\eta$ , the normalized mean of the  $f(c_i^T) - f(c_i^F)$ . The crucial insight is that if we have a diverse set of counterfactual pairs, then in the hyperplane orthogonal to  $a$ , the fluctuations of  $f(c_i^T) - f(c_i^F)$  get averaged and should disappear when taking the mean. Therefore,  $\eta$  will be parallel to  $a$ , which is the optimal steering vector. This intuition can be generalized to non-atomic concepts, as we will describe in Appendix D.3. In summary, our framework enables progress on mechanistic interpretability [66], an emerging field which aims to understand LLM mechanisms.

Next, we propose a modification of this technique as follows. First, we use our insights to arrive at the intuition that any weighted mean of the counterfactual differences should roughly have the same steering behavior. And moreover, we choose the weights dynamically, with weights being similarities of the test sentence to the sentences from the training set. This enables similar sentences in the training set to have a higher vote in the choice of steering vector. With this new idea, we use as steering vector

$$\eta(x) = \sum_i \langle \lambda(x), \lambda(c_i^F) \rangle (f(c_i^T) - f(c_i^F))$$

where  $\lambda$  is a sentence embedding (such as SBERT [79]). We detail this further in Appendix D.3. Finally, we also show how to implement this efficiently using *steering matrices* instead of steering vectors.

Table 2: Comparison of steering vectors for LLM alignment

Technique	$\alpha$	Acc (%)	CE loss	KL div.
Baseline	-	25.7	2.16	0.0
Random direction	20	25.8	2.19	0.02
CCS direction	5	26.2	2.21	0.06
ITI: Probe weight dir.	15	27.0	2.21	0.06
ITI: Mass mean shift	20	28.8	2.41	0.27
Steering matrices (ours)	15	<b>29.5</b>	2.61	0.41

In Table 2, we show the results of our experiments with steering matrices. We use the open-source large language model LLaMA [97] with 7 billion parameters and the sentence transformer SBERT [79] for the sentence embedding. We report the accuracy of the multiple-choice track of TruthfulQA [51] over 3 random seeds and also the Cross-Entropy Loss and KL divergence of the model pre- and post-intervention. Higher accuracy is better and lower CE loss, and KL divergence indicate that the original model has not

<sup>1</sup>Note that the simplified version is stated here. In ITI,  $\eta$  is not actually a constant vector because it's an autoregressive model. However a constant vector is added to each layer during inference time (which is sufficient for relaying intuition in this section). See Appendix D for details.

been significantly modified. Here, the baselines are the unmodified model, random direction intervention, Contrast-Consistent Search (CCS) direction [12] and two different direction choices using vanilla ITI; and 2-fold cross validation is used. Additional technical details are deferred to Appendix D.

We see that the multiple-choice accuracy improved, showcasing the potential of our steering matrices technique which is novel in the field of LLM alignment to the best of our knowledge. Note that this is meant to be a proof of concept and not meant to be a comprehensive study of this specific technique. For exploratory purposes, we outline potential modifications to our technique in Appendix D which could potentially improve the performance, both in terms of accuracy as well as in terms of invasiveness. These form an exciting direction for a more comprehensive study of our proposed ideas, which we leave for future work.

## 6 Conclusion

In this work, we exploited ideas from both the causal representation learning and the foundations model literature to study the problem of learning human-interpretable concepts from data. We proposed a formal definition of concepts and studied under what conditions they can be provably recovered, suggesting what representations foundation models learn. Finally, synthetic experiments and large language model alignment experiments showcase the utility of our ideas.

While we have made initial progress in unifying these disparate fields, this direction holds a lot of promise. We believe there's a lot of upside in utilizing ideas from identifiable representation learning in order to theoretically analyze what representations foundation models learn, as it will help us understand, explain, and improve their capabilities.

**Acknowledgments** We acknowledge the support of AFRL and DARPA via FA8750-23-2-1015, ONR via N00014-23-1-2368, NSF via IIS-1909816, IIS-1955532, IIS-1956330, and NIH R01GM140467. We also acknowledge the support of the Tübingen AI Center and the Robert H. Topel Faculty Research Fund at the University of Chicago Booth School of Business.

## References

- [1] K. Ahuja, D. Mahajan, Y. Wang, and Y. Bengio. Interventional causal representation learning. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org, 2023.
- [2] K. Ahuja, A. Mansouri, and Y. Wang. Multi-domain causal representation learning via weak distributional invariances. *arXiv preprint arXiv:2310.02854*, 2023.
- [3] C. Allen and T. Hospedales. Analogies explained: Towards understanding word embeddings. In *International Conference on Machine Learning*, pages 223–231. PMLR, 2019.
- [4] S. Arora, Y. Li, Y. Liang, T. Ma, and A. Risteski. A latent variable model approach to pmi-based word embeddings. *Transactions of the Association for Computational Linguistics*, 4:385–399, 2016.
- [5] Y. Bai, A. Jones, K. Ndousse, A. Askell, A. Chen, N. DasSarma, D. Drain, S. Fort, D. Ganguli, T. Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- [6] Y. Bai, S. Kadavath, S. Kundu, A. Askell, J. Kernion, A. Jones, A. Chen, A. Goldie, A. Mirhosseini, C. McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- [7] D. Bau, B. Zhou, A. Khosla, A. Oliva, and A. Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6541–6549, 2017.
- [8] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [9] J. Brehmer, P. De Haan, P. Lippe, and T. S. Cohen. Weakly supervised causal representation learning. *Advances in Neural Information Processing Systems*, 35:38319–38331, 2022.
- [10] S. Buchholz, M. Besserve, and B. Schölkopf. Function classes for identifiable nonlinear independent

component analysis. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=DpKaP-PY8bK>.

- [11] S. Buchholz, G. Rajendran, E. Rosenfeld, B. Aragam, B. Schölkopf, and P. Ravikumar. Learning linear causal representations from interventions under general nonlinear mixing. *arXiv preprint arXiv:2306.02235*, 2023.
- [12] C. Burns, H. Ye, D. Klein, and J. Steinhardt. Discovering latent knowledge in language models without supervision. *arXiv preprint arXiv:2212.03827*, 2022.
- [13] D. V. Carvalho, E. M. Pereira, and J. S. Cardoso. Machine learning interpretability: A survey on methods and metrics. *Electronics*, 8(8):832, 2019.
- [14] W.-L. Chiang, Z. Li, Z. Lin, Y. Sheng, Z. Wu, H. Zhang, L. Zheng, S. Zhuang, Y. Zhuang, J. E. Gonzalez, I. Stoica, and E. P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna/>.
- [15] P. Comon. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994.
- [16] A. Conneau, G. Kruszewski, G. Lample, L. Barrault, and M. Baroni. What you can cram into a single vector: Probing sentence embeddings for linguistic properties. *arXiv preprint arXiv:1805.01070*, 2018.
- [17] J. Cui, W. Huang, Y. Wang, and Y. Wang. Aggnce: Asymptotically identifiable contrastive learning. In *NeurIPS Workshop*, 2022.
- [18] A. D’Amour, K. Heller, D. Moldovan, B. Adlam, B. Alipanahi, A. Beutel, C. Chen, J. Deaton, J. Eisenstein, M. D. Hoffman, et al. Underspecification presents challenges for credibility in modern machine learning. *The Journal of Machine Learning Research*, 23(1):10237–10297, 2022.
- [19] N. Dilokthanakul, P. A. Mediano, M. Garnelo, M. C. Lee, H. Salimbeni, K. Arulkumaran, and M. Shanahan. Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint arXiv:1611.02648*, 2016.
- [20] M. Drton, B. Sturmfels, and S. Sullivant. *Lectures on Algebraic Statistics*, volume 39 of *Oberwolfach Seminars*. Springer, 2009. doi: 10.1007/978-3-7643-8905-5.
- [21] N. Elhage, N. Nanda, C. Olsson, T. Henighan, N. Joseph, B. Mann, A. Askell, Y. Bai, A. Chen, T. Conerly, et al. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1, 2021.
- [22] N. Elhage, T. Hume, C. Olsson, N. Schiefer, T. Henighan, S. Kravec, Z. Hatfield-Dodds, R. Lasenby, D. Drain, C. Chen, R. Grosse, S. McCandlish, J. Kaplan, D. Amodei, M. Wattenberg, and C. Olah. Toy models of superposition. *Transformer Circuits Thread*, 2022. [https://transformer-circuits.pub/2022/toy\\_model/index.html](https://transformer-circuits.pub/2022/toy_model/index.html).
- [23] J. Engel, M. Hoffman, and A. Roberts. Latent constraints: Learning to generate conditionally from unconditional generative models. *arXiv preprint arXiv:1711.05772*, 2017.
- [24] K. Ethayarajh, D. Duvenaud, and G. Hirst. Towards understanding linear word analogies. *arXiv preprint arXiv:1810.04882*, 2018.
- [25] F. Falck, H. Zhang, M. Willetts, G. Nicholson, C. Yau, and C. C. Holmes. Multi-facet clustering variational autoencoders. *Advances in Neural Information Processing Systems*, 34, 2021.
- [26] D. Ganguli, L. Lovitt, J. Kernion, A. Askell, Y. Bai, S. Kadavath, B. Mann, E. Perez, N. Schiefer, K. Ndousse, et al. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*, 2022.
- [27] A. Gittens, D. Achlioptas, and M. W. Mahoney. Skip-gram- zipf+ uniform= vector additivity. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 69–76, 2017.
- [28] L. Greselle, J. Von Kügelgen, V. Stimper, B. Schölkopf, and M. Besserve. Independent mechanism analysis, a new concept? *Advances in Neural Information Processing Systems*, 34, 2021.
- [29] S. Gupta, S. Jegelka, D. Lopez-Paz, and K. Ahuja. Context is environment. *arXiv e-prints*, pages arXiv–2309, 2023.

[30] W. Gurnee, N. Nanda, M. Pauly, K. Harvey, D. Troitskii, and D. Bertsimas. Finding neurons in a haystack: Case studies with sparse probing. *arXiv preprint arXiv:2305.01610*, 2023.

[31] E. Hernandez, B. Z. Li, and J. Andreas. Measuring and manipulating knowledge representations in language models. *arXiv preprint arXiv:2304.00740*, 2023.

[32] A. Hyvärinen and H. Morioka. Unsupervised feature extraction by time-contrastive learning and nonlinear ica. *Advances in neural information processing systems*, 29, 2016.

[33] A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural networks*, 13(4-5):411–430, 2000.

[34] A. Hyvärinen and P. Pajunen. Nonlinear independent component analysis: Existence and uniqueness results. *Neural networks*, 12(3):429–439, 1999.

[35] A. Hyvärinen, J. Karhunen, and E. Oja. Independent component analysis. *Studies in informatics and control*, 11(2):205–207, 2002.

[36] A. Hyvärinen, H. Sasaki, and R. Turner. Nonlinear ica using auxiliary variables and generalized contrastive learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 859–868. PMLR, 2019.

[37] A. Hyvärinen, I. Khemakhem, and R. Monti. Identifiability of latent-variable and structural-equation models: from linear to nonlinear. *arXiv preprint arXiv:2302.02672*, 2023.

[38] Y. Jiang and B. Aragam. Learning latent causal graphs with unknown interventions. In *Advances in Neural Information Processing Systems*, 2023.

[39] Y. Jiang, B. Aragam, and V. Veitch. Uncovering meanings of embeddings via partial orthogonality. *Advances in Neural Information Processing Systems*, 2023.

[40] Y. Jiang, G. Rajendran, P. Ravikumar, B. Aragam, and V. Veitch. On the origins of linear representations in large language models. *arXiv preprint*, 2024.

[41] I. Khemakhem, D. Kingma, R. Monti, and A. Hyvärinen. Variational autoencoders and nonlinear ica: A unifying framework. In *International Conference on Artificial Intelligence and Statistics*, pages 2207–2217. PMLR, 2020.

[42] I. Khemakhem, R. Monti, D. Kingma, and A. Hyvärinen. Ice-beem: Identifiable conditional energy-based deep models based on nonlinear ica. *Advances in Neural Information Processing Systems*, 33:12768–12778, 2020.

[43] B. Kim, M. Wattenberg, J. Gilmer, C. Cai, J. Wexler, F. Viegas, et al. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning*, pages 2668–2677. PMLR, 2018.

[44] B. Kivva, G. Rajendran, P. Ravikumar, and B. Aragam. Learning latent causal graphs via mixture oracles. *Advances in Neural Information Processing Systems*, 34:18087–18101, 2021.

[45] B. Kivva, G. Rajendran, P. Ravikumar, and B. Aragam. Identifiability of deep generative models without auxiliary information. *Advances in Neural Information Processing Systems*, 35:15687–15701, 2022.

[46] L. Kong, S. Xie, W. Yao, Y. Zheng, G. Chen, P. Stojanov, V. Akinwande, and K. Zhang. Partial identifiability for domain adaptation. *arXiv preprint arXiv:2306.06510*, 2023.

[47] S. Lachapelle, P. Rodríguez, Y. Sharma, K. Everett, R. L. Priol, A. Lacoste, and S. Lacoste-Julien. Disentanglement via mechanism sparsity regularization: A new principle for nonlinear ICA. In B. Schölkopf, C. Uhler, and K. Zhang, editors, *1st Conference on Causal Learning and Reasoning, CLeaR 2022, Sequoia Conference Center, Eureka, CA, USA, 11-13 April, 2022*, volume 177 of *Proceedings of Machine Learning Research*, pages 428–484. PMLR, 2022. URL <https://proceedings.mlr.press/v177/lachapelle22a.html>.

[48] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

[49] O. Levy and Y. Goldberg. Neural word embedding as implicit matrix factorization. *Advances in neural information processing systems*, 27, 2014.

[50] K. Li, A. K. Hopkins, D. Bau, F. Viégas, H. Pfister, and M. Wattenberg. Emergent world representations: Exploring a sequence model trained on a synthetic task. *arXiv preprint arXiv:2210.13382*, 2022.

2022.

- [51] K. Li, O. Patel, F. Viégas, H. Pfister, and M. Wattenberg. Inference-time intervention: Eliciting truthful answers from a language model. *arXiv preprint arXiv:2306.03341*, 2023.
- [52] S. Li, B. Hooi, and G. H. Lee. Identifying through flows for recovering latent representations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=Sk10UpEYvB>.
- [53] S. Lin, J. Hilton, and O. Evans. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*, 2021.
- [54] P. Lippe, S. Magliacane, S. Löwe, Y. M. Asano, T. Cohen, and E. Gavves. Biscuit: Causal representation learning from binary interactions. *arXiv preprint arXiv:2306.09643*, 2023.
- [55] Y. Liu, Z. Zhang, D. Gong, M. Gong, B. Huang, A. v. d. Hengel, K. Zhang, and J. Q. Shi. Identifying weight-variant latent causal models. *arXiv preprint arXiv:2208.14153*, 2022.
- [56] F. Locatello, S. Bauer, M. Lucic, G. Raetsch, S. Gelly, B. Schölkopf, and O. Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, pages 4114–4124. PMLR, 2019.
- [57] I. Loshchilov and F. Hutter. SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=Skq89Scxx>.
- [58] C. Lu, Y. Wu, J. M. Hernández-Lobato, and B. Schölkopf. Invariant causal representation learning for out-of-distribution generalization. In *International Conference on Learning Representations*, 2021.
- [59] T. McGrath, A. Kapishnikov, N. Tomašev, A. Pearce, M. Wattenberg, D. Hassabis, B. Kim, U. Paquet, and V. Kramnik. Acquisition of chess knowledge in alphazero. *Proceedings of the National Academy of Sciences*, 119(47):e2206625119, 2022.
- [60] K. Meng, D. Bau, A. Andonian, and Y. Belinkov. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372, 2022.
- [61] T. Mikolov, W.-t. Yih, and G. Zweig. Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 conference of the north american chapter of the association for computational linguistics: Human language technologies*, pages 746–751, 2013.
- [62] G. E. Moran, D. Sridhar, Y. Wang, and D. Blei. Identifiable deep generative models via sparse decoding. *Transactions on Machine Learning Research*, 2022.
- [63] L. Moschella, V. Maiorca, M. Fumero, A. Norelli, F. Locatello, and E. Rodola. Relative representations enable zero-shot latent space communication. *arXiv preprint arXiv:2209.15430*, 2022.
- [64] R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
- [65] N. Nanda, A. Lee, and M. Wattenberg. Emergent linear representations in world models of self-supervised sequence models. *arXiv preprint arXiv:2309.00941*, 2023.
- [66] C. Olah. Mechanistic interpretability, variables, and the importance of interpretable bases. <https://transformer-circuits.pub/2022/mech-interp-essay/index.html>, 2022.
- [67] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [68] K. Park, Y. J. Choe, and V. Veitch. The linear representation hypothesis and the geometry of large language models. *arXiv preprint arXiv:2311.03658*, 2023.
- [69] J. Pearl. *Causality*. Cambridge university press, 2009.
- [70] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.

[71] J. Peters, D. Janzing, and B. Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.

[72] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

[73] A. Radford, R. Jozefowicz, and I. Sutskever. Learning to generate reviews and discovering sentiment. *arXiv preprint arXiv:1704.01444*, 2017.

[74] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.

[75] R. Rafailov, A. Sharma, E. Mitchell, S. Ermon, C. D. Manning, and C. Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023.

[76] M. Raghu, J. Gilmer, J. Yosinski, and J. Sohl-Dickstein. Svcca: Singular vector canonical correlation analysis for deep understanding and improvement. *stat*, 1050:19, 2017.

[77] G. Rajendran, B. Kivva, M. Gao, and B. Aragam. Structure learning in polynomial time: Greedy algorithms, bregman information, and exponential families. *Advances in Neural Information Processing Systems*, 34:18660–18672, 2021.

[78] G. Rajendran, P. Reizinger, W. Brendel, and P. Ravikumar. An interventional perspective on identifiability in gaussian lti systems with independent component analysis. *arXiv preprint arXiv:2311.18048*, 2023.

[79] N. Reimers and I. Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. URL <https://arxiv.org/abs/1908.10084>.

[80] N. Rimsky, N. Gabrieli, J. Schulz, M. Tong, E. Hubinger, and A. M. Turner. Steering llama 2 via contrastive activation addition. *arXiv preprint arXiv:2312.06681*, 2023.

[81] B. Schölkopf and J. von Kügelgen. From statistical to causal learning. In *Proceedings of the International Congress of Mathematicians (ICM)*. EMS Press, July 2022.

[82] B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. arXiv:2102.11107.

[83] L. Schut, N. Tomasev, T. McGrath, D. Hassabis, U. Paquet, and B. Kim. Bridging the human-ai knowledge gap: Concept discovery and transfer in alphazero. *arXiv preprint arXiv:2310.16410*, 2023.

[84] Y. Seonwoo, S. Park, D. Kim, and A. Oh. Additive compositionality of word vectors. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 387–396, 2019.

[85] X. Shen, F. Liu, H. Dong, Q. Lian, Z. Chen, and T. Zhang. Weakly supervised disentangled generative causal representation learning. *Journal of Machine Learning Research*, 23:1–55, 2022.

[86] K. Shuster, S. Poff, M. Chen, D. Kiela, and J. Weston. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*, 2021.

[87] P. Sorrenson, C. Rother, and U. Köthe. Disentanglement by nonlinear ica with general incompressible-flow networks (gin). *arXiv preprint arXiv:2001.04872*, 2020.

[88] P. Spirtes, C. N. Glymour, and R. Scheines. *Causation, prediction, and search*. MIT press, 2000.

[89] C. Squires and C. Uhler. Causal structure learning: a combinatorial perspective. *Foundations of Computational Mathematics*, pages 1–35, 2022.

[90] C. Squires, A. Seigal, S. S. Bhate, and C. Uhler. Linear causal disentanglement via interventions. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, editors, *International Conference on Machine Learning, ICML 2023, 23–29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 32540–32560. PMLR, 2023. URL <https://proceedings.mlr.press/v202/squires23a.html>.

[91] N. Subramani, N. Suresh, and M. E. Peters. Extracting latent steering vectors from pretrained language models. *arXiv preprint arXiv:2205.05124*, 2022.

[92] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

[93] D. Talon, P. Lippe, S. James, A. Del Bue, and S. Magliacane. Towards the reusability and compositionality of causal representations. In *Causal Representation Learning Workshop at NeurIPS 2023*, 2023.

[94] R. Taori, I. Gulrajani, T. Zhang, Y. Dubois, X. Li, C. Guestrin, P. Liang, and T. B. Hashimoto. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models. https://crfm. stanford. edu/2023/03/13/alpaca. html*, 3(6):7, 2023.

[95] I. Tenney, D. Das, and E. Pavlick. Bert rediscovers the classical nlp pipeline. *arXiv preprint arXiv:1905.05950*, 2019.

[96] C. Tigges, O. J. Hollinsworth, A. Geiger, and N. Nanda. Linear representations of sentiment in large language models. *arXiv preprint arXiv:2310.15154*, 2023.

[97] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

[98] M. Trager, P. Perera, L. Zancato, A. Achille, P. Bhatia, and S. Soatto. Linear spaces of meanings: compositional structures in vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15395–15404, 2023.

[99] A. Turner, L. Thiergart, D. Udell, G. Leech, U. Mini, and M. MacDiarmid. Activation addition: Steering language models without optimization. *arXiv preprint arXiv:2308.10248*, 2023.

[100] B. Varici, K. Shanmugam, P. Sattigeri, and A. Tajer. Intervention target estimation in the presence of latent variables. In *Uncertainty in Artificial Intelligence*, pages 2013–2023. PMLR, 2022.

[101] B. Varici, E. Acarturk, K. Shanmugam, A. Kumar, and A. Tajer. Score-based causal representation learning with interventions. *arXiv preprint arXiv:2301.08230*, 2023.

[102] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[103] J. Von Kügelgen, Y. Sharma, L. Gresele, W. Brendel, B. Schölkopf, M. Besserve, and F. Locatello. Self-supervised learning with data augmentations provably isolates content from style. *Advances in Neural Information Processing Systems*, 34, 2021.

[104] J. von Kügelgen, M. Besserve, W. Liang, L. Gresele, A. Kekić, E. Bareinboim, D. M. Blei, and B. Schölkopf. Nonparametric identifiability of causal representations from unknown interventions. In *Advances in Neural Information Processing Systems*, 2023.

[105] K. Wang, A. Variengien, A. Conmy, B. Shlegeris, and J. Steinhardt. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small. *arXiv preprint arXiv:2211.00593*, 2022.

[106] Y. Wang, D. Blei, and J. P. Cunningham. Posterior collapse and latent variable non-identifiability. *Advances in Neural Information Processing Systems*, 34:5443–5455, 2021.

[107] Z. Wang, L. Gui, J. Negrea, and V. Veitch. Concept algebra for score-based conditional model. In *ICML 2023 Workshop on Structured Probabilistic Inference \{\backslash\mathcal{E}\} Generative Modeling*, 2023.

[108] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.

[109] M. Willets and B. Paige. I don't need **u**: Identifiable non-linear ica without side information. *arXiv preprint arXiv:2106.05238*, 2021.

[110] M. Yang, F. Liu, Z. Chen, X. Shen, J. Hao, and J. Wang. Causalvae: Disentangled representation learning via neural structural causal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9593–9602, June 2021.

[111] F. Zhang and N. Nanda. Towards best practices of activation patching in language models: Metrics and methods. *arXiv preprint arXiv:2309.16042*, 2023.

[112] J. Zhang, C. Squires, K. Greenewald, A. Srivastava, K. Shanmugam, and C. Uhler. Identifiability guarantees for causal disentanglement from soft interventions. *arXiv preprint arXiv:2307.06250*,

2023.

- [113] Y. Zhang, Y. Du, B. Huang, Z. Wang, J. Wang, M. Fang, and M. Pechenizkiy. Interpretable reward redistribution in reinforcement learning: A causal approach. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [114] Y. Zheng, I. Ng, and K. Zhang. On the identifiability of nonlinear ica: Sparsity and beyond. *Advances in Neural Information Processing Systems*, 35:16411–16422, 2022.
- [115] R. S. Zimmermann, Y. Sharma, S. Schneider, M. Bethge, and W. Brendel. Contrastive learning inverts the data generating process. In *International Conference on Machine Learning*, pages 12979–12990. PMLR, 2021.
- [116] A. Zou, L. Phan, S. Chen, J. Campbell, P. Guo, R. Ren, A. Pan, X. Yin, M. Mazeika, A.-K. Dombrowski, et al. Representation engineering: A top-down approach to ai transparency. *arXiv preprint arXiv:2310.01405*, 2023.

## A Proofs of the main results

In this appendix we provide the proofs of our results, in particular the proof of our main result, Theorem 1. However, as a warm-up we first start in Appendix A.1 with a proof of the simpler result that can be shown based on the iVAE approach. In Appendix A.2 we prove Theorem 1 and in Appendix A.3 we prove the additional lemmas that appear in the paper.

### A.1 Proof of identifiability with $2n + 1$ environments

As a warm-up and to provide a connection to earlier results we show here how to obtain identifiability by adapting the iVAE framework to our context. Indeed, our mathematical setting is related to the setting used in [41] in the sense that the environments are generated by modulation with certain exponential families. Therefore, we can essentially apply their proof techniques to prove identifiability (with some modifications), albeit this requires the suboptimal number of  $2m + 1$  environments (there are two sufficient statistics for the Gaussian distribution).

**Theorem 2.** *Suppose data satisfies Assumption 1, 2, and 3 and the environment statistics matrix  $\Lambda$  has rank  $2n$ . Assume we know the number of atoms  $n$ . Then identifiability in the sense of Definition 4 holds.*

We remark that the rank condition can only be satisfied for  $2n + 1$  environments (observational distribution and  $2n$  concept conditional distributions. For this theorem the assumption that the filtering distribution is always the same is not necessary. Instead we could consider variances  $(\sigma_k^e)^2$  depending on environment  $e$  and row  $k$ , i.e., the filtering distribution  $q_{(\sigma_k^e)^2}$  is Gaussian with varying variance. The generalization of the environment-concept matrix  $M \in \mathbb{R}^{m \times n}$  is given by

$$M_{ei} = \begin{cases} \frac{1}{(\sigma_k^e)^2} & \text{if } i \in S^e \text{ and row } k \text{ of } A^e \text{ is } a_i \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

Similarly the generalization of the environment-valuation matrix  $B \in \mathbb{R}^{m \times n}$  is given by

$$B_{ei} = \begin{cases} \frac{b_k^e}{(\sigma_k^e)^2} & \text{if } i \in S^e \text{ and row } k \text{ of } A^e \text{ is } a_i, \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

We now prove Theorem 2. We use essentially the same ideas as in the proof of Theorem 1 in Khemakhem et al. [41] (followed by the same reasoning as in Sorrenson et al. [87], Kivva et al. [45] but since our concepts are not axis aligned and we only extract some information about the mixing we give a complete proof).

*Proof of Theorem 2.* Suppose there are 2 sets of parameters that generate the same data  $X^0, X^1, \dots, X^m$ . Denote by  $\tilde{\cdot}$  the latter set of parameters, e.g.,  $X^e$  is distributed as  $\tilde{f}(\tilde{Z}^e)$  where  $\tilde{Z}^e \in \mathbb{R}^{\tilde{d}_z}$  corresponds to the concept class  $\tilde{C}^e$  with distribution  $\tilde{Z}^e \sim \tilde{p}^e$  and the same distribution is generated by  $f(Z^e)$  where  $f$  and  $\tilde{f}$  are injective and differentiable. Let  $\mathcal{C} = \{a_1, \dots, a_n\}$  be the set of atomic concepts in the first setting and let  $\tilde{\mathcal{C}} = \{\tilde{a}_1, \dots, \tilde{a}_n\}$  be the set of atomic concepts in the second setting (here we use that  $n$  is assumed to be known). We also consider the transition function  $\varphi = \tilde{f}^{-1}f$  and in the following we always write  $\tilde{Z} = \varphi(Z)$ . The equality  $f(Z^e) \stackrel{\mathcal{D}}{=} X^e \stackrel{\mathcal{D}}{=} \tilde{f}(\tilde{Z}^e)$  implies  $\varphi(Z^e) \stackrel{\mathcal{D}}{=} \tilde{Z}^e$ . This implies that for all environments  $e$

$$p^e(Z) = |\det J_{\varphi^{-1}}| \cdot \tilde{p}^e(\tilde{Z}) \quad (11)$$

Taking the logarithm and subtracting this for some  $e = 1, \dots, m$  from the base distribution we obtain

$$\ln(p(Z)) - \ln(p^e(Z)) = \ln(\tilde{p}(\tilde{Z})) - \ln(\tilde{p}^e(\tilde{Z})). \quad (12)$$

Using the definition (2) we can rewrite for some constants  $c_e$  and  $c'_e$

$$\begin{aligned} \ln(p(Z)) - \ln(p^e(Z)) &= \sum_{k=1}^{\dim(C_e)} \frac{(A^e Z^e - b^e)_k^2}{2(\sigma_k^e)^2} - c'_e \\ &= \sum_{i=1}^n \left( \frac{1}{2} M_{ei} \langle a_i, Z^e \rangle^2 - B_{ei} \langle a_i, Z^e \rangle \right) - c_e. \end{aligned} \quad (13)$$

Here we used the environment-concept matrix and the environment-valuation matrix in the second step which were defined in (7) and (8) (in (9) and (10) for varying variance). We define the vector  $\mathbf{p}(Z)$  with components  $\mathbf{p}_e(Z) = \ln(p(Z)) - \ln(p^e(Z))$ . Then we find the relation

$$\mathbf{p}(Z) = \frac{1}{2}M \begin{pmatrix} \langle a_1, Z \rangle^2 \\ \vdots \\ \langle a_n, Z \rangle^2 \end{pmatrix} - B \begin{pmatrix} \langle a_1, Z \rangle \\ \vdots \\ \langle a_n, Z \rangle \end{pmatrix}. \quad (14)$$

Together with (12) we conclude that

$$\frac{1}{2}M \begin{pmatrix} \langle a_1, Z \rangle^2 \\ \vdots \\ \langle a_n, Z \rangle^2 \end{pmatrix} - B \begin{pmatrix} \langle a_1, Z \rangle \\ \vdots \\ \langle a_n, Z \rangle \end{pmatrix} = \frac{1}{2}\widetilde{M} \begin{pmatrix} \langle \tilde{a}_1, \tilde{Z} \rangle^2 \\ \vdots \\ \langle \tilde{a}_n, \tilde{Z} \rangle^2 \end{pmatrix} - \widetilde{B} \begin{pmatrix} \langle \tilde{a}_1, \tilde{Z} \rangle \\ \vdots \\ \langle \tilde{a}_n, \tilde{Z} \rangle \end{pmatrix} \quad (15)$$

Since by assumption  $\widetilde{\Lambda} = (\widetilde{M}, \widetilde{B}) \in \mathbb{R}^{m \times 2n}$  has rank  $2n$  there is a vector  $v$  such that  $v^\top \widetilde{M} = 0$  and  $v^\top \widetilde{B} = -\mathbf{e}_i$  ( $\mathbf{e}_i \in \mathbb{R}^{d_z}$  denotes the  $i$ -th standard basis vector). Thus we find that

$$\langle \tilde{a}_i, \tilde{Z} \rangle = \frac{1}{2}v^\top M \begin{pmatrix} \langle a_1, Z \rangle^2 \\ \vdots \\ \langle a_n, Z \rangle^2 \end{pmatrix} - v^\top B \begin{pmatrix} \langle a_1, Z \rangle \\ \vdots \\ \langle a_n, Z \rangle \end{pmatrix}. \quad (16)$$

In other words  $\langle \tilde{a}_i, \tilde{Z} \rangle$  can be expressed as a quadratic polynomial in  $Z$ . We apply the same reasoning for  $\langle \tilde{a}_i, \tilde{Z} \rangle^2$ , i.e., pick a vector  $v'$  such that  $\frac{1}{2}v'^\top \widetilde{M} = \mathbf{e}_i$  and  $v'^\top \widetilde{B} = 0$  to obtain a relation

$$\langle \tilde{a}_i, \tilde{Z} \rangle^2 = \sum_j \eta_j \langle a_j, Z \rangle^2 + \ell(Z) \quad (17)$$

for some coefficients  $\eta_j$  and some affine function  $\ell$  of  $Z$ . The following reasoning is now the same as in Kivva et al. [45], Sorrenson et al. [87]. We thus find that  $\langle \tilde{a}_i, \tilde{Z} \rangle$  and its square can be written as polynomials of degree at most 2 in  $Z$ . This implies that in fact  $\langle \tilde{a}_i, \tilde{Z} \rangle$  is an affine function of  $Z$  (otherwise its square would be a quartic polynomial), i.e., we can write

$$\langle \tilde{a}_i, \tilde{Z} \rangle = \sum_j \lambda_j \langle a_j, Z \rangle + C_i = \langle \sum_j \lambda_j a_j, Z \rangle + C_i. \quad (18)$$

Equating the square of this relation with (17) and taking the gradient with respect to  $Z$  (as a polynomial the function is differentiable) we find

$$2 \sum_j \eta_j a_j \langle a_j, Z \rangle + w = 2 \sum_j \lambda_j a_j \langle \sum_j \lambda_j a_j, Z \rangle + w' \quad (19)$$

for two vectors  $w$  and  $w'$ . The equality (for  $Z = 0$ ) implies  $w = w'$ . Now linear independence of  $a_j$  implies that for each  $r$

$$\eta_r a_r = \lambda_r \sum_j \lambda_j a_j. \quad (20)$$

Applying linear independence again we conclude that either  $\lambda_r = 0$  or  $\lambda_j = 0$  for all  $j \neq r$ . This implies that there is at most one  $r$  such that  $\lambda_r \neq 0$ . The relation (18) and the bijectivity of  $\varphi$  implies that there is exactly one  $r(i)$  such that  $\lambda_{r(i)} \neq 0$  and therefore

$$\langle \tilde{a}_i, \tilde{Z} \rangle = \lambda_{r(i)} \langle a_{r(i)}, Z \rangle + C_i. \quad (21)$$

Applying the same argument in the reverse direction we conclude that there is a permutation  $\pi \in S_n$  such that

$$\langle \tilde{a}_{\pi(i)}, \tilde{Z} \rangle = \lambda_i \langle a_i, Z \rangle + C_i. \quad (22)$$

By linear independence we can find an invertible linear map  $T$  such that

$$\tilde{a}_{\pi(i)}^\top = a_i^\top T^{-1} \quad (23)$$

(i.e.,  $T^\top \tilde{a}_{\pi(i)} = a_i$ ) and a vector  $w \in \mathbb{R}^{d_z}$  (the  $a_i$  are linearly independent) such that

$$\langle \tilde{a}_{\pi(i)}, \tilde{Z} \rangle = \lambda_i(\langle a_i, Z \rangle + \langle a_i, w \rangle). \quad (24)$$

In particular the relations (5) and (6) hold. Now it is straightforward to see that if  $i \in S^e$ , i.e.,  $a_i$  is a row of  $A^e$  then  $\tilde{a}_{\pi(i)}$  is a row of  $\tilde{A}^e$  and vice versa. Indeed, this follows from (15) for environment  $e$  together with (24) and linear independence of the atoms. Therefore we conclude from (23) that there is a permutation  $P^e$  such that

$$\tilde{A}^e = P^e A^e T^{-1}. \quad (25)$$

Moreover, (24) then implies setting  $Z = f^{-1}(x)$ ,  $\tilde{Z} = \tilde{f}^{-1}(x)$

$$\tilde{A}^e \tilde{f}^{-1}(x) = \Lambda^e P^e A^e (f^{-1}(x) + w) \quad (26)$$

holds for the same permutation matrix  $P^e$  and a diagonal matrix  $\Lambda^e$  whose diagonal entries can be related to (24). Let us assume now that row  $k$  of  $A^e$  is  $a_i$  and row  $k'$  of  $\tilde{A}^e$  is  $\tilde{a}_{\pi(i)}$ . Now we consider the subspace  $H \subset \mathbb{R}^{d_z}$  containing all  $Z$  such that  $\langle Z, a_j \rangle = 0$  for  $j \neq i$ . Via (24) this implies that  $\langle \tilde{a}_j, \tilde{Z} \rangle$  is constant for  $j \neq \pi(i)$ . Then we conclude from (15) that for  $Z \in H$

$$\frac{(\langle a_i, Z \rangle - b_k^e)^2}{2(\sigma_k^e)^2} = \frac{(\langle \tilde{a}_{\pi(i)}, \tilde{Z} \rangle - \tilde{b}_{k'}^e)^2}{2(\tilde{\sigma}_{k'}^e)^2} + c_k^e \quad (27)$$

for some constant  $c_k^e$ . Using (24) this implies that

$$\frac{(\langle a_i, Z \rangle - b_k^e)^2}{2(\sigma_k^e)^2} = \frac{(\lambda_i(\langle a_i, Z \rangle + \langle a_i, w \rangle) - \tilde{b}_{k'}^e)^2}{2(\tilde{\sigma}_{k'}^e)^2} + c_k^e. \quad (28)$$

Comparing the quadratic term and the linear term (note that  $\langle a_i, Z \rangle$  can take any value on  $H$ ) we find

$$\frac{1}{2(\sigma_k^e)^2} = \frac{\lambda_i^2}{2(\tilde{\sigma}_{k'}^e)^2} \quad (29)$$

$$-\frac{b_k^e}{2(\sigma_k^e)^2} = -\frac{\lambda_i \tilde{b}_{k'}^e - \lambda_i^2 \langle a_i, w \rangle}{2(\tilde{\sigma}_{k'}^e)^2} \quad (30)$$

Combining the equation we obtain

$$\tilde{b}_{k'}^e = \lambda_i(b_k^e - \langle a_i, w \rangle) \quad (31)$$

This implies then the relation

$$\tilde{b} = \Lambda^e P^e (b + A^e w). \quad (32)$$

□

## A.2 Proof of Theorem 1

In this section we prove our main Theorem 1. The proof is structured in several steps: First we remove the symmetries of the representation and derive the key relations underlying the proof. Then we show that we can identify the environment-concept matrix  $M$  and then also the valuations collected in  $B$ . Once this is done we can complete the proof. We will need the following lemma to conclude the proof.

**Lemma 3.** *The relations (3) and (6) in Definition 4 define an equivalence relation of representations if we assume that the underlying atoms form a linearly independent set.*

The proof of this lemma can be found in Appendix A.3.

**Remark 3.** *Without the assumption on the underlying atoms the lemma is not true. In this case a slightly different scaling must be chosen (e.g.,  $(\Lambda^e)^{-1} \tilde{b}^e = \Lambda^e P^e b^e - P^e A^e w$  instead of  $\tilde{b}^e = \Lambda^e P^e (b^e - A^e w)$ ). Since our results address the case of atoms we used the simpler definition in the main paper.*

We can allow slightly more general filtering distributions where  $q$  is Gaussian with variance  $\sigma_i^2$  if we filter on concept  $i$ , i.e., the variance needs to be constant for different environments and the same atom but might depend on the atom. The proof will cover this case, the simple case stated in the main paper is obtained by setting  $\sigma_i^2 = \sigma^2$ . Some steps of the proof (e.g., the expressions for the difference of the log-densities) agree with the proof of Theorem 2. To keep the proof self contained we repeat a few equations.

*Proof of Theorem 1.* We proceed in several steps.

**Step 1: Reduction to standard form.** Let us first transform every possible data representation into a standard form. Recall that we have the set of atomic concepts  $\mathcal{C} = \{a_1, \dots, a_n\}$ . Recall that we defined the environment-concept matrix  $M \in \mathbb{R}^{m \times n}$  in (7) and note that the natural generalisation reads

$$M_{ei} = \begin{cases} \frac{1}{\sigma_i^2} & \text{if } a_i \text{ is a row of } A^e, \\ 0 & \text{otherwise.} \end{cases} \quad (33)$$

We say that concept  $a_n$  is conditioned on the environment  $e$ . Note that the nonzero entries of row  $e$  of  $M$  encode the set  $S^e$ . To pass from  $A^e$  to its rows  $a_i$  we assume that the  $e$ -th row of  $A^e$  is  $a_{i_j^e}$ , i.e.,  $a_{i_j^e} = (A^e)^\top e_j$ . Recall also consider the environment-valuation matrix  $B$  which is given by

$$B_{ei} = \begin{cases} \frac{b_k^e}{\sigma_i^2} & \text{if } a_i \text{ is row } k \text{ of } A^e, \\ 0 & \text{otherwise.} \end{cases} \quad (34)$$

Denoting by  $q_{\sigma^2}$  the centered Gaussian distribution with variance  $\sigma^2$  we find in environment  $e$

$$\begin{aligned} \ln(p(Z)) - \ln(p^e(Z)) &= - \sum_{k=1}^{\dim(C_e)} \ln q_{(\sigma_k^e)^2}((A^e Z^e - b^e)_k) = \sum_{k=1}^{\dim(C_e)} \frac{(A^e Z^e - b^e)_k^2}{2(\sigma_k^e)^2} - c'_e \\ &= \sum_{i=1}^n \frac{1}{2} M_{ei} \langle a_i, Z^e \rangle^2 - B_{ei} \langle a_i, Z^e \rangle - c_e. \end{aligned} \quad (35)$$

Now we consider an invertible linear map  $T : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_z}$  such that  $T^{-\top} a_i = e_i$  for all  $1 \leq i \leq n$ . Such a map exists because we assume that the  $a_i$  are linearly independent. Moreover, we consider a shift vector  $\lambda \in \mathbb{R}^{d_z}$  with  $\lambda_i = 0$  for  $i > n$  which we fix later. We define  $\Sigma \in \mathbb{R}^{d_z \times d_z}$  to be the diagonal matrix with entries  $\Sigma_{ii} = \sigma_i$  for  $1 \leq i \leq n$  and  $\Sigma_{ii} = 1$  for  $i > n$ . Now we consider the linear map  $L(z) = \Sigma^{-1} T z - \lambda$  and a new representation given by

$$\bar{z} = L(z), \quad \bar{f} = f \circ L^{-1}, \quad \bar{\mathcal{C}} = \{e_1, \dots, e_n\}, \quad \bar{\sigma}_i = 1, \quad \bar{A}^e = A^e T^{-1}, \quad \bar{p}(\bar{z}) = p(L^{-1} \bar{z}) |\det T^{-1}|. \quad (36)$$

We also define

$$\bar{b}_k^e = \frac{b_k^e}{\sigma_i} - \lambda_i \quad \text{if row } k \text{ of } A^e \text{ is } a_i. \quad (37)$$

Define  $\bar{M}$  and  $\bar{B}$  in terms of  $\bar{A}^e$ ,  $\bar{b}^e$  and  $\bar{\sigma}_i^2$  as before. We remark that all entries of  $\bar{M}$  are either 0 or 1 and note that

$$\bar{M} = M \text{Diag}(\sigma_1^2, \dots, \sigma_n^2) \quad (38)$$

$$\bar{B} = B \text{Diag}(\sigma_1^{-1}, \dots, \sigma_n^{-1}) - M \text{Diag}(\lambda_1, \dots, \lambda_n). \quad (39)$$

We claim that this model generates the same observations as the original model. By definition  $L_* p = \bar{p}$  (as mentioned before, we slightly abuse notation and here refer to the distributions). Next, we calculate for any  $\delta$

$$\begin{aligned} -2 \ln q_1(\langle e_i, L(z) \rangle - \delta) &= (\langle e_i, L(z) \rangle - \delta)^2 \\ &= (\langle e_i, \Sigma T z - \lambda \rangle - \delta)^2 \\ &= (\sigma_i^{-1} \langle T^\top e_i, z \rangle - \lambda_i - \delta)^2 \\ &= \frac{(\langle a_i, z \rangle - \sigma_i \lambda_i - \sigma_i \delta)^2}{\sigma_i^2} \\ &= -2 \ln q_{\sigma_i^2}(\langle a_i, z \rangle - \sigma_i \lambda_i - \sigma_i \delta). \end{aligned} \quad (40)$$

Using this for  $\delta = \bar{b}_k^e$  and some  $k$  such that row  $k$  of  $A^e$  is  $a_i$  we find

$$-2 \ln q_1(\langle e_i, L(z) \rangle - \bar{b}_k^e) = -2 \ln q_{\sigma_i^2}(\langle a_i, z \rangle - \sigma_i \lambda_i - \sigma_i \bar{b}_k^e) = -2 \ln q_{\sigma_i^2}(\langle a_i, z \rangle - b_k^e). \quad (41)$$

This then implies that for  $\tilde{z} = L(z)$

$$\prod_k q_1((\tilde{A}^e \tilde{z} - \tilde{b}^e)_k) \propto \prod_k q_{\sigma_k^e}((A^e z - b^e)_k). \quad (42)$$

Combining this with the definition (2) and the definition  $\bar{p}(\tilde{z}) = p(L^{-1}\tilde{z})|\det T^{-1}|$  we find that for  $\bar{z} = L(z)$

$$\bar{p}^e(\bar{z}) \propto p^e(z) \quad (43)$$

and thus  $\bar{f}(\bar{Z}^e) \stackrel{D}{=} f(Z^e)$ . Moreover, one directly sees that the two representations are also equivalent in the sense of Definition 4. We now fix the vector  $\lambda$  such that each row of  $\bar{B}$  has mean zero. Finally, by changing the sign of  $\tilde{z}_i$  we can in addition assume that for every  $i$  the first non-zero  $\bar{B}_{ei}$  is positive. Finally we remark that Assumption 4 is still satisfied for  $\bar{M}$  and  $\bar{B}$ . Indeed,  $w^\top M = 0$  implies  $w^\top \bar{M} = 0$  by (38). But then  $w^\top \bar{B} = w^\top B \text{Diag}(\sigma_1^{-1}, \dots, \sigma_n^{-1})$  by (39) which has all entries different from zero if this holds for  $w^\top B$ . In the following we will therefore always assume that the representation satisfies the properties of the  $\bar{Z}$  variables and we remove the modifier in the following. The plan is now to show that  $M$  and  $B$  can be identified up to permutations of the rows (under the fixed normalization we derived in this step) and then show that every two representations with the same  $M$  and  $B$  can be identified.

**Step 2: The key identity** Let us here restate the key identity based on the difference of the log-densities. As is common in identifiability results for multi-environment data with general mixing we consider the difference in log densities. Consider

$$\begin{aligned} \ln p^0(z) - \ln p^e(z) &= \sum_{i=1}^n \frac{1}{2} M_{ei} \langle \mathbf{e}_i, z \rangle^2 - B_{ei} \langle \mathbf{e}_i, z \rangle - c'_e \\ &= \sum_{i=1}^n \frac{1}{2} M_{ei} z_i^2 - B_{ei} z_i - c'_e \end{aligned} \quad (44)$$

for some constant  $c'_e$ . Those functions will play a crucial role in the following and we will denote

$$g^e(z) = \ln p^0(z) - \ln p^e(z) \quad (45)$$

Note that since the log-density changes only by the Jacobian for pushforward measures we find that

$$g^e(z) = \ln p^0(z) - \ln p^e(z) = \ln p_X^0(f(z)) - \ln p_X^e(f(z)) = G^e(f(z)) = G^e(x). \quad (46)$$

Note that the functions  $G^e(x)$  can be estimated from the distributions of  $X^e$ . We remark  $X$  might be supported on a submanifold if  $d_z$  and  $d_x$  do not agree making the definition of the density subtle. But we can just consider any chart locally and consider the density of the pushforward with respect to the Lebesgue measure. The resulting difference expressed in  $G^e$  will be independent of the chart as the determinant cancels thus  $G^e$  is a well defined function. The relation

$$g^e(z) = G^e(f(z)) = G^e(x) \quad (47)$$

will be crucial in the following because it shows that properties of  $g^e$  are closely linked to the identifiable functions  $G^e$ .

**Step 3: Identifiability of environment-concept matrix** Let us now show that we can identify which concepts are contained in which environment (up to relabeling of the concepts). Recall that  $S^e = \{i \in [n] : a_i \text{ is a row of } A^e\}$  and we similarly define  $S_T = \bigcup_{e \in T} S^e$  for all subsets  $T \subset [m]$ . The main observation is that we can identify  $|S_T| = |\bigcup_{e \in T} S^e|$  for all subsets  $T \subset [m]$ . To show this we consider the set

$$I_T = \operatorname{argmin}_z \sum_{e \in T} g^e(z). \quad (48)$$

Note that the function  $g^e$  are convex functions, and they can be decomposed as sums of functions in  $z_i$ , i.e., for some functions  $h_i^T$

$$\sum_{e \in T} g^e(z) = \sum_{i=1}^n h_i^T(z_i). \quad (49)$$

Now if  $i \in S_T$  then  $i \in S^e$  for some  $e$  and thus  $M_{ei} \neq 0$  for the  $e$  and  $h_i^T$  is the sum of quadratic function in  $x_i$  which as a strictly convex function has a unique minimum  $z_i^T$ . On the other hand, if  $i \notin S_T$  then  $i \notin S^e$  for  $e \in T$  and thus  $M_{ei} = 0$  for all  $e \in T$  and  $h_i^T(z_i) = 0$ . Thus we conclude that

$$I_T = \{z \in \mathbb{R}^{d_z} : z_i = z_i^T \text{ for } i \in S_T\}. \quad (50)$$

This is an affine subspace of dimension  $d_z - |S_T|$ . The relations  $G^e(f(z)) = g^e(z)$  imply that

$$f(I_T) = \operatorname{argmin}_x \sum_{e \in T} G^e(x). \quad (51)$$

Note that  $G^e(x)$  is identifiable from the datasets  $X^e$  and thus the submanifold (by assumption on  $f$ )  $f(I_T)$  is identifiable and by finding its dimension we obtain  $d_z - |S_T|$ . Since  $d_z$  is the dimension of the data manifold  $f(X)$  we can indeed identify  $|S_T|$  for all  $T \subset [m]$ . In particular, the total number of atomic concepts  $n = |S_{[m]}|$  is identifiable (assuming that all atomic concepts are filtered upon at least once). Now, it is a standard result that we can identify the matrix  $M$  up to permutation of the atomic concepts. Indeed, we can argue by induction in  $m$  to show this. For  $m = 1$  we just have  $|S^1|$  atomic concepts appearing in environment 1 and  $n - |S^1|$  concepts not appearing. For the induction step  $m \rightarrow m + 1$  we consider the sizes  $|S_{T \cup \{m+1\}}|$  for  $T \subset [m]$ . Applying the induction hypothesis we can complete  $M_{ei}$  for all columns such that  $M_{m+1,i} = 1$ . Similarly, we can consider the sizes  $|S_T| - |S_{T \cup \{m+1\}}|$  to identify the matrix  $M$  for concepts not used in environment  $m + 1$ .

Thus, we can and will assume after permuting the atomic concepts that  $M$  is some fixed matrix.

**Step 4: Identifiability of concept valuations** Next, we show that we can also identify the matrix  $B$ . We do this column by column, i.e., for one atomic concept after another. Assume we consider atomic concept  $i$ . Then we consider the set  $T_i = \{e : M_{ei} = 0\}$  of concepts that not filter on atomic concept  $i$ . By Assumption 5 there is for every  $i' \neq i$  an environment  $e$  such that  $i'$  is filtered on, i.e.,  $M_{ei'} \neq 0$ . This implies  $S_{T_i} = [n] \setminus \{i\}$ . Then we consider as in (50) the set  $I_{T_i}$  given by

$$I_{T_i} = \{z \in \mathbb{R}^{d_z} : z_{i'} = z_{i'}^{T_i} \text{ for } i' \in [n] \setminus \{i\}\}. \quad (52)$$

Note that all  $z_{i'}$  for  $i \neq i'$  are constant on  $I_{T_i}$ . Thus we find for any environment  $e$  such that  $i \in S^e$ .

$$\begin{aligned} g^e(z) &= \sum_{j=1}^n \frac{1}{2} M_{ej} z_j^2 - B_{ej} z_j - c'_e \\ &= \sum_{j \neq i} \frac{1}{2} M_{ej} z_j^2 - B_{ej} z_j - c'_e + \frac{1}{2} z_i^2 - B_{ei} z_i \\ &= c_{T_i, e} + \frac{1}{2} z_i^2 - B_{ei} z_i \end{aligned} \quad (53)$$

on  $I_{T_i}$  for some constant  $c_{T_i}$ .

Now we consider two concepts  $e_1 \neq e_2$  such that atomic concept  $i$  is contained in these two environments. Then we consider the set

$$I_{T_i}^{e_1} = \operatorname{argmin}_{z \in I_{T_i}} g^{e_1}(z) = \{z \in \mathbb{R}^{d_z} : z_{i'} = z_{i'}^{T_i} \text{ for } i' \in [n] \setminus \{i\}, z_i = B_{e_1 i}\}. \quad (54)$$

Note that in the second equality we used that  $g^{e_1}(z)$  depends on  $z_i$  through  $z_i^2/2 - B_{e_1 i} z_i$  so it is minimized at  $B_{e_1 i}$ . Now we find using (53)

$$\begin{aligned} \min_{z \in I_{T_i}^{e_1}} g^{e_2}(z) - \min_{I_{T_i}} g^{e_2}(z) &= \min_{z \in I_{T_i}^{e_1}} c_{T_i, e_2} + \frac{1}{2} z_i^2 - B_{e_2 i} z_i - \min_{I_{T_i}} \left( c_{T_i, e_2} + \frac{1}{2} z_i^2 - B_{e_2 i} z_i \right) \\ &= c_{T_i, e_2} + \frac{1}{2} B_{e_1 i}^2 - B_{e_1 i} B_{e_2 i} - \left( c_{T_i, e_2} + \frac{1}{2} B_{e_2 i}^2 - B_{e_2 i}^2 \right) \\ &= \frac{(B_{e_1 i} - B_{e_2 i})^2}{2}. \end{aligned} \quad (55)$$

As before, this quantity is identifiable from observations because  $f(T_i)$  can be identified and we can minimize  $G^{e_2}(x)$  over  $f(T_i)$ .

This allows us to identify  $B_{e_1i} - B_{e_2i}$  up to a sign. However, we can evaluate this expression over all pairs  $e_1$  and  $e_2$  and pick the one with the maximal difference. Then all remaining values  $B_{ei}$  for  $e$  such that  $i$  is filtered on in  $e$  must satisfy  $B_{ei} \in [B_{e_1i}, B_{e_2i}]$ . Together with identifiability of  $|B_{ei} - B_{e_1i}|$  this allows us to identify all  $B_{ei}$  up to one sign indeterminacy and a constant shift. However, in the first step we ensured that  $\sum_e B_{ei} = 0$  for all  $i$  which determines the shift and the sign is fixed by our choice of making the first non-zero entry positive. Thus, we can assume that our two representations have the same  $M$  and  $B$ .

**Step 5: Identifiability of concepts** We are now ready to prove our identifiability result.

Assume we have two representations  $Z^e, f, p$  and  $\tilde{Z}^e, \tilde{f}, \tilde{p}$  such that the corresponding environment-concept and environment-valuation matrices agree, i.e.,  $M = \tilde{M}$  and  $B = \tilde{B}$ . We consider the transition function  $\varphi = \tilde{f}^{-1} \circ f$  which is by assumption differentiable. What we want to show is that  $\varphi(z)_i = z_i$  for all  $z \in \mathbb{R}^{d_z}$  and  $1 \leq i \leq n$ . We now decompose  $z = (z^c, z^o)$  into the concept part and the orthogonal part. We fix  $z^o \in \mathbb{R}^{d_z-n}$  and define the function  $\iota^o(z^c) = (z^c, z^o)$ , the projection  $\pi^c((z^c, z^o)) = z^c$ , and  $\varphi^o : \mathbb{R}^n \rightarrow \mathbb{R}^n$  given by  $\varphi^o(z^c)_i = \varphi(\iota^o(z^c))_i = \varphi((z^c, z^o))_i$ . Note that  $\varphi^o$  is differentiable but not necessarily injective. Let us denote by  $\mathbf{g} : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^m$  the function with coordinates  $\mathbf{g}_e = g^e$  and similarly we define  $\mathbf{G} : M \rightarrow \mathbb{R}^d$ . Identifiability will be based on the crucial relation

$$\mathbf{g}(\iota^o(z^c)) = \mathbf{G}(f(\iota^o(z^c))) = \mathbf{G}(\tilde{f}(\varphi^o(z^c))) = \mathbf{g}(\varphi^o(z^c)). \quad (56)$$

Here we used in the last step that  $g^e$  is defined in terms of  $M$  and  $B$  and thus agrees for both representations. Note that  $\mathbf{g}$  is just a quadratic function. Differentiating we obtain

$$D_i g^e(z) = M_{ei} z_i - B_{ei}. \quad (57)$$

Concisely this can be written as

$$D\mathbf{g} = M \text{Diag}(z_1, \dots, z_n) - B. \quad (58)$$

Differentiating (56) we find

$$M \text{Diag}(z_1, \dots, z_n) - B = (M \text{Diag}(\tilde{z}_1, \dots, \tilde{z}_n) - B) D\varphi^o(z^c). \quad (59)$$

Let  $v$  be a vector as in Assumption 4. Denote by  $M^+ \in \mathbb{R}^{n \times m}$  the pseudoinverse of  $M$  which has rank  $n$  because  $M$  has. We consider the matrix  $\widetilde{M}^+ \in \mathbb{R}^{n+1 \times m}$  given by

$$\widetilde{M}^+ = \begin{pmatrix} M^+ \\ v^\top \end{pmatrix} \quad (60)$$

Let us multiply the relation (59) by  $\widetilde{M}^+$  and find that

$$\begin{pmatrix} z_1 & & 0 \\ & \ddots & \\ 0 & & z_n \\ 0 & \dots & 0 \end{pmatrix} - \widetilde{M}^+ B = \left( \begin{pmatrix} \tilde{z}_1 & & 0 \\ & \ddots & \\ 0 & & \tilde{z}_n \\ 0 & \dots & 0 \end{pmatrix} - \widetilde{M}^+ B \right) D\varphi^o(z^c) \quad (61)$$

Note that the first  $n$  rows of the left hand side are  $\text{Diag}(z_1, \dots, z_n) - M^+ B$ . This matrix is invertible for almost all values of  $z^c = (z_1, \dots, z_n)^\top$  because its determinant is a non-zero polynomial (the coefficient of the term  $z_1 \dots z_n$  is 1) which vanishes only on a set of measure zero. Outside of this set the left hand side of has rank  $n$ . Then the equality (61) implies that also the right hand side has rank  $n$  and thus  $D\varphi^o(z^c)$  has rank  $n$  and thus is invertible. For  $z^c$  outside of this set there is up to scaling a unique vector  $w \neq 0$  (depending on  $z_1, \dots, z_n$ ) such that

$$w^\top \left( \begin{pmatrix} z_1 & & 0 \\ & \ddots & \\ 0 & & z_n \\ 0 & \dots & 0 \end{pmatrix} - \widetilde{M}^+ B \right) = 0 \quad (62)$$

From (61) we conclude using the invertibility of  $D\varphi^o(z^c)$  that

$$w^\top \left( \begin{pmatrix} \tilde{z}_1 & & 0 \\ & \ddots & \\ 0 & & \tilde{z}_n \\ 0 & \dots & 0 \end{pmatrix} - \widetilde{M^+} B \right) = 0. \quad (63)$$

Next, we claim that for almost all values of  $z^c$  the vector  $w$  has all entries different from 0 (this property is invariant under rescaling). Actually we need this only for entries 1 to  $n$  but the case  $n+1$  is a bit simpler so we show it first. We show this by proving that for each entry  $w_i$  there is only a null set of  $z^c$  such that  $w_i = 0$ . Let  $w = (w', 0)$  for some  $w' \in \mathbb{R}^n$  and  $w' \neq 0$ , i.e.,  $w_{n+1} = 0$ . Then

$$0 = w^\top \left( \begin{pmatrix} z_1 & & 0 \\ & \ddots & \\ 0 & & z_n \\ 0 & \dots & 0 \end{pmatrix} - \widetilde{M^+} B \right) = w'^\top (\text{Diag}(z_1, \dots, z_n) - M^+ B) \quad (64)$$

But this implies that  $\text{Diag}(z_1, \dots, z_n) - M^+ B$  has non-trivial kernel, i.e., does not have full rank and we have seen above that this happens only for a subset of measure 0 of all  $z^c$ . Next we show that the same is true if  $w_1 = 0$ . Decompose  $0 \neq w = (0, w')$ . Then we find

$$0 = w^\top \left( \begin{pmatrix} z_1 & & 0 \\ & \ddots & \\ 0 & & z_n \\ 0 & \dots & 0 \end{pmatrix} - \widetilde{M^+} B \right) = w'^\top \left( \begin{pmatrix} 0 & z_2 & 0 & 0 \\ \dots & \ddots & & \\ 0 & & z_n & \\ 0 & \dots & \dots & 0 \end{pmatrix} - (\widetilde{M^+} B)_{2:(n+1)} \right) \quad (65)$$

Thus we conclude that the matrix on the right hand side is not invertible. Its determinant is a polynomial in  $z_2, \dots, z_n$  and its highest degree term is  $\pm z_2 \cdot \dots \cdot z_n \cdot (\widetilde{M^+} B)_{(n+1),1}$ . By definition of  $\widetilde{M^+} B$  we find  $(\widetilde{M^+} B)_{(n+1),1} = (v^\top B)_1 \neq 0$  by Assumption 4 (recall that we showed invariance of the assumption under the transformation of  $M$  and  $B$ ). We find that the determinant is a non-zero polynomial and the set of its zeros is a set of measure 0 of all  $z_2, \dots, z_n$  but since it does not depend on  $z_1$  this holds true for almost all  $z^c$ . The same reasoning for  $i = 2, \dots, n$  implies that for every  $i$  the set of  $z^c$  such that  $w_i = 0$  is a set of measure zero. We have therefore shown that for almost all  $z^c$  the rank of the left hand side of (61) is  $n$  and the corresponding vector  $w \neq 0$  has all entries different from zero. Subtracting (62) and (63) we obtain

$$0 = w^\top \left( \begin{pmatrix} z_1 & & 0 \\ & \ddots & \\ 0 & & z_n \\ 0 & \dots & 0 \end{pmatrix} - w^\top \begin{pmatrix} \tilde{z}_1 & & 0 \\ & \ddots & \\ 0 & & \tilde{z}_n \\ 0 & \dots & 0 \end{pmatrix} \right) = (w_1(z_1 - \tilde{z}_1), \dots, w_n(z_n - \tilde{z}_n), 0). \quad (66)$$

Now  $w_i \neq 0$  implies  $z_i = \tilde{z}_i$ . We conclude that for almost all  $z^c$  the relation  $\varphi^o(z^c) = z^c$  holds. By continuity this implies that the relation actually holds everywhere. We conclude that  $\pi^c \tilde{f}^{-1} f((z^c, z^o)) = z^c$  for a fixed  $z^o$  but since  $z^o$  was arbitrary the relation holds for all  $z^o$  and all  $z^c$ . Thus we conclude that for  $1 \leq i \leq n$

$$\langle \mathbf{e}_i, \tilde{f}^{-1}(x) \rangle = \langle \mathbf{e}_i, \varphi(f^{-1}(x)) \rangle = \langle \mathbf{e}_i, f^{-1}(x) \rangle \quad (67)$$

holds. This implies that those two representations satisfy (3) and (4) (with  $P^e = \Lambda^e = \text{Id}$  and  $T = \text{Id}$ ). But since this relation is an equivalence relation in our setting by Lemma 3 and since we showed equivalence to a representation in standard form in the first step we conclude that also any two representations are related through (3) and (4) thus finishing the proof.  $\square$

### A.3 Remaining proofs

Here we prove the remaining auxiliary results.

*Proof of Lemma 1.* Since  $M \in \mathbb{R}^{m \times n}$  has rank  $n$  and  $m = n + 1$  there is exactly one vector  $v \in \mathbb{R}^m$  such that  $v^\top M = 0$  and  $v \neq 0$ . We claim that this vector has all entries different from zero. Indeed suppose  $v_m = 0$  which then implies  $v_{1:(m-1)}^\top M_{1:(m-1)} = 0$ . But by assumption every  $n \times n$  submatrix of  $M$  is invertible (this is equivalent to the rows being linearly independent) so we conclude that  $v_{1:(m-1)} = 0$  which is a contradiction to  $v \neq 0$ . The same reasoning applies to every entry. Note that the assumption on  $M$  implies that every column has at least one non-zero entry, i.e., every column of  $B$  has one entry sampled from a continuous distribution. But then the probability that  $v$  is orthogonal to a column is zero because this is a codimension 1 hyperplane of all valuations of this row (since all entries of  $v$  are non-zero).  $\square$

*Proof of Lemma 3.* Reflexivity is obvious, just pick  $T = \text{Id}$ ,  $w = 0$ ,  $\Lambda^e = P^e = \text{Id}_{\dim(C^e)}$ . To show symmetry we first consider the atoms. Let  $\tilde{T} = T^{-1}$  and  $\tilde{\pi} = \pi^{-1}$ . Then

$$a_{\tilde{\pi}(i)}^\top = a_{\pi^{-1}(i)}^\top T^{-1} T = \tilde{a}_{\pi \circ \pi^{-1}(i)} \tilde{T}^{-1} = \tilde{a}_i \tilde{T}^{-1}. \quad (68)$$

Let  $\tilde{w}$  be a vector such that for all  $1 \leq i \leq n$

$$\langle a_i, w \rangle = -\frac{1}{\lambda_i} \langle \tilde{a}_{\pi(i)}, \tilde{w} \rangle. \quad (69)$$

Such a vector exists by linear independence of  $\tilde{a}_i$ . Let  $\tilde{\lambda}_i = \lambda_{\tilde{\pi}(i)}^{-1}$ . Then we find that the relation (6), namely

$$\langle \tilde{a}_{\pi(i)}, \tilde{f}^{-1}(x) \rangle = \lambda_i (\langle a_i, f^{-1}(x) \rangle + \langle a_i, w \rangle) \quad (70)$$

implies

$$\begin{aligned} \langle a_{\tilde{\pi}(i)}, f^{-1}(x) \rangle &= \frac{1}{\lambda_{\tilde{\pi}(i)}} \langle \tilde{a}_{\pi \circ \tilde{\pi}(i)}, \tilde{f}^{-1}(x) \rangle - \langle a_{\tilde{\pi}(i)}, w \rangle = \frac{1}{\lambda_{\tilde{\pi}(i)}} \langle \tilde{a}_i, \tilde{f}^{-1}(x) \rangle + \frac{1}{\lambda_{\tilde{\pi}(i)}} \langle \tilde{a}_{\pi \circ \tilde{\pi}(i)}, \tilde{w} \rangle \\ &= \tilde{\lambda}_i (\langle \tilde{a}_i, \tilde{f}^{-1}(x) \rangle + \langle \tilde{a}_i, \tilde{w} \rangle). \end{aligned} \quad (71)$$

It remains to be shown that this lifts to the concepts  $C^e$ . We first note that the relation (6) together with (69) and (3) implies that

$$\Lambda^e P^e A^e w = -\tilde{A}^e \tilde{w}. \quad (72)$$

Let  $\tilde{P}^e = (P^e)^{-1}$  and  $\tilde{\Lambda}^e = (P^e)^{-1}(\Lambda^e)^{-1}P^e$ . Then (3) combined with the previous display implies

$$A^e f^{-1}(x) = (P^e)^{-1}(\Lambda^e)^{-1} \tilde{A}^e \tilde{f}^{-1}(x) - A^e w = \tilde{\Lambda}^e \tilde{P}^e \tilde{A}^e \tilde{f}^{-1}(x) + (P^e)^{-1}(\Lambda^e)^{-1} \tilde{A} \tilde{w} = \tilde{\Lambda}^e \tilde{P}^e \tilde{A}^e (\tilde{f}^{-1}(x) + \tilde{w}). \quad (73)$$

The relation

$$A^e = \tilde{P}^e \tilde{A}^e \tilde{T}^{-1} \quad (74)$$

is a direct consequence of the definitions of  $\tilde{P}^e$  and  $\tilde{T}$  and (4) and the relation

$$b^e = \tilde{\Lambda}^e \tilde{P}^e (\tilde{b}^e - \tilde{A}^e w) \quad (75)$$

follows exactly as in (73). The proof of transitivity is similar (first establish the relations on the atomic concepts then lift it to  $C^e$ ).  $\square$

## B Comparison to Causal Representation Learning

In this appendix we describe causal representation learning and discuss the similarities and differences between the viewpoint taken in this paper and the standard setting in causal representation learning.

Causal Representation Learning (CRL) [82, 81] is a modern machine learning field that aims to learn representations of data that correspond to the true generative process. More precisely, if we assume that data  $X$  is generated as  $X = f(Z)$  where  $Z$  are latent causal factors and  $f$  is some arbitrary nonlinearity, the goal is to learn  $f$  as well as the distribution of  $Z$ . To make this more aligned with reality, the latent

variables  $Z$  are assumed to have causal relationships among them. For instance,  $Z_1, Z_2$  could correspond to shape and size of an object respectively and  $f$  could correspond to the rendering of an image of the object. Then, using datasets of images, we wish to learn latent variables that correspond to shape and size respectively, along with the image rendering map. CRL incorporates ideas from the field of causality [88, 69, 71, 77, 89] into the field of latent variable models and is a generalization of nonlinear independent component analysis [15, 33, 35] and disentangled representation learning [8, 71, 48]. The field has seen a surge of advances in the last few years, e.g., [41, 44, 25, 55, 47, 10, 62, 115, 28, 78, 101, 38, 37, 93] and featured several workshops on Causal Representation Learning, among others at UAI 2022, CLeaR 2023 and NeurIPS 2023. As motivated in Schölkopf et al. [82], CRL enables many desiderata such as robustness, out of distribution generalization, and in addition enables planning and alignment. CRL has been successfully applied to many domains, already showing fascinating progress in genomics [112, 90] and holds great promise for other domains such as vision or text (see references above).

At the same time, research into foundation models for text, images and other fields has burgeoned in recent years. Significant progress has been made to enable robustness, prediction and generalization among other desiderata. Much of this research has been largely experimental and has not strictly followed the paradigms set forth in causal representation learning. We therefore endeavor to answer this discrepancy in this work with our unified framework.

To describe the intuition, we will re-emphasize the goal of CRL below, because our work critically departs from this goal.

*Causal representation learning aims to find the true underlying generative factors for a data distribution.*

The issue here is that attempts towards CRL often conflate the notion of *causal generative factors* with *human interpretable factors*. That is, it's tempting to imagine that generative factors for a dataset are natural causes that are interpretable to human beings. However, there is no reason that this should indeed be the case.

In our work, we take significant inspiration from the framework of causal representation learning and present a slightly relaxed paradigm that is weaker, but more general and also importantly, aligns better with many high-performance foundation models in the literature. We now describe the setup of CRL more formally in Appendix B.1. Then, in Appendix B.2, we discuss conceptual differences between causal representation learning and our framework.

## B.1 Formal setup

We assume that we observe data  $X \in \mathbb{R}^{d_x}$  with the generative model  $X = f(Z)$  where  $Z \in \mathbb{R}^{d_z}$  are the latent variables and  $f$  is a deterministic mixing function. The dataset  $X$  is sampled from a distribution  $p$  and the goal is to recover the mixing function  $f$  as well as the distributions of the underlying latent variables  $Z_1, \dots, Z_{d_z}$ . To this end, this problem is over-parameterized since multiple pairs of  $Z$  and  $f$  could fit the dataset apriori, so the field of representation learning makes various assumptions to learn this model *identifiably*. Here, identifiability is the notion that a unique set of parameters fit the model (up to trivial transformations). This makes the problem well-defined and feasible (however it could still be a hard problem to solve in practice). Below, we informally summarize two classes of prior works that enable such identifiability guarantees.

1. Disentangled representation learning: In this setting, we assume that the distributions of  $Z_1, \dots, Z_{d_z}$  are jointly independent. Different studies constrain the distribution of the variables  $Z_1, \dots, Z_{d_z}$ , e.g., each  $Z_i$  is independently sampled from  $N(0, 1)$ . This is also the setting studied in nonlinear independent component analysis [15, 33].
2. Causal Representation Learning: This setting is more general than the one above where we relax the independence assumption on the  $Z_i$ , instead we assume that they have causal relationships among them. For instance, they could satisfy a linear structural causal model with Gaussian noise, i.e.,  $Z = AZ + \epsilon, \epsilon \sim N(0, I)$  where  $A$  encodes a weighted directed acyclic graph. As stated, this setting is more general than the above, since having no causal relationships (i.e.,  $A = 0$ ) essentially indicates joint independence.

As explained earlier, in both these domains, a critical notion is that of identifiability [41, 18, 106], which posits that the given dataset(s) are diverse enough for the modeling assumptions, in order to ensure that a unique set of parameters fit the data.

It's folklore that the disentangled representation learning model is not identifiable if all  $Z_i$  are Gaussian [34, 56]. However, under appropriate assumptions, e.g., distributional, sparsity or observed side-information, the model becomes identifiable, see e.g., Khemakhem et al. [41], Hyvarinen and Morioka [32], Brehmer et al. [9], Shen et al. [85], Lachapelle et al. [47], Moran et al. [62], Zheng et al. [114], Kivva et al. [45], Buchholz et al. [10], Zimmermann et al. [115], Gresele et al. [28], Rajendran et al. [78]. In addition, various works have proposed methods to learn them (with modest success) [25, 109, 19, 110, 52, 17, 10, 11].

## B.2 Conceptual differences

In this section, we highlight the conceptual differences between causal representation learning and our framework.

**Are causal generative concepts necessarily interpretable?** While it's tempting to imagine the true causal generative factors as being naturally human interpretable concepts (e.g., intuitive abstractions such as shape, size or color of an object), there's no obvious reason why this should be the case. The perspective that we take in this work is that the number of true generative factors could be prohibitively large so that attempting to learn them is infeasible, whereas the number of desired human-interpretable concepts is much smaller. Moreover, we constantly come up with new concepts of interest since human-interpretable concepts are constantly evolving, e.g., the concept of mobile phones was not existent 100 years ago, but is a valid concept to learn now. Therefore, as opposed to working with a rigid model as in causal representation learning, we take the approach of working with a dynamic representation learning model.

**Number of environments needed** When the ground truth generative process has ambient latent dimension  $d_z$ , for causal representation learning to be feasible, we usually require  $d_z$  environments or datasets. For instance, in the iVAE setting above [41], we require  $d_z k + 1 \geq d_z + 1$  environments. This is indeed necessary, as counterexamples show. However, it's not clear what the value of  $d_z$  is for complex datasets, and it could potentially be prohibitively large.

But the question remains, do we need to learn the entire generative model for solving downstream tasks? Along these lines, there is a tremendous research effort attempting to relax such requirements by imposing various inductive or domain biases and by building a theory of partial identifiability [45, 54, 46]. This is for good reason, since even though it would be ideal to learn the full ground truth generative model, it may be too ambitious or prohibitively large and moreover it maybe not be necessary for most downstream tasks we care about, therefore it suffices to learn what's necessary. On this note, the relevant task of learning only a subset of the generative latent variables is not easy as the latent variables interact in potentially complicated ways and research in this front is non-trivial and limited progress has been made.

In this work however, we show that if we only wish to learn  $m \ll d_z$  concepts, it suffices to have  $O(m)$  environments instead of  $O(d_z)$  environments. Therefore, our results can be viewed as a result on partial identifiability with sublinear number of environments.

**Multi-node interventions** Multi-node interventions are an exciting area of study in CRL, since they are a natural extension of existing works and are more useful for modeling various real-life datasets where it can be hard to control precisely one factor of variation. This is easily incorporated in our setting by utilizing non-atomic concepts, since each non-atomic concept is a collection of vectors corresponding to atomic concepts and can be modified simultaneously by changing the valuation.

**Conditional vs. interventional data** In this work we focus on conditional data and identification of concept values, while most of the recent identifiability results focus on interventional data and identification of the causal structure [89, 100, 11, 38, 103]. In general conditional data is more frequently available. Conditional data can be obtained by selection through filtering, e.g., patients that are admitted to different hospitals based on the severity of their condition or by the availability of label information as in the CLIP setting [74]. Thus conditional data can be obtained by observing the system in different conditions. On the other hand interventional data requires manipulation of the system which is more difficult to obtain in general. Moreover, many of the theoretical results require perfect interventions which can be achieved by randomized control trials but are otherwise a very strong assumption. We note that in typical applications we do not think of concepts as being causal variables that are connected by a graph. On the other hand, our framework can accomodate the case where concepts correspond to causal

variables. Our identifiability result does not allow us to learn the causal structure. However, identifying the causal variables reduces the problem from causal representation learning to causal inference (with observational data) which is still a difficult problem but nevertheless substantially simpler than CRL.

## C Alternate definitions of concept conditional measure

In this section, we present alternate feasible definitions for data distributions than the one we introduced in Section 3.2. While we went with the definition most suited for practice, these alternate definitions are also justifiable in different scenarios and are exciting avenues for further study.

We want to essentially define a concept  $C$  via a conditional measure  $p_C$  where the concept  $C$  is identified with an affine subspace  $C = \{Z \in \mathbb{R}^{d_z} : A^C Z = b^C\}$  for some  $A^C \in \mathbb{R}^{k \times d_z}$ ,  $b^C \in \mathbb{R}^k$ . We consider the shifted parallel linear subspace  $C_0 = \{Z : A^C Z = 0\}$  and the orthogonal splitting  $\mathbb{R}^{d_z} = C_0 \oplus V$ . Suppose we have a distribution  $q_V$  on the space  $V$  which will typically be a Gaussian centred around  $v^C \in V$  which is the unique solution of  $A^C v^C = b^C$ . In addition we have a base distribution  $p$  on  $\mathbb{R}^{d_z}$ . We will assume that all distributions have a smooth density so that conditional probabilities are pointwise well defined. There are at least three ways to create the context conditional measure  $p_C$ .

1. The first option is to enforce that the distribution of the  $V$  marginal  $p_C(v) = \int_{C_0} p_C(v, c) dc$  exactly matches  $q_V(v)$  while the in-plane distribution  $p_C(c|v = v_0) \propto p_C(c, v_0)$  remains invariant, i.e., equals  $p(c|v = v_0)$ . Under this condition, there is a unique measure  $p_C$  given by

$$p_C(c, v) \propto q_V(v) \frac{p(c, v)}{\int_{C_0} p(c', v) dc'}.$$

In other words, to get  $(c, v)$  we sample  $v \sim q_V$  and then  $c \sim p(c|v)$  according to the conditional distribution.

2. The second option is to again enforce the  $V$  marginal but instead of keeping the in plane distribution we average over the  $V$  space. Then we obtain

$$p_C(c, v) \propto q_V(v) \int_V p(c, v') dv'.$$

This corresponds (vaguely) to a  $\text{do}(v)$  operation from causal inference, i.e., we sample according to  $p(v, c)$  and then do a random intervention on  $v$  with target distribution  $q_V$ .

3. The third option is to take a Bayesian standpoint. Then we view  $p$  as a prior and  $q_V$  as the context dependent acceptance probability, i.e., we sample by  $p$  and then accept with probability  $q_V$ . Then we find

$$p_C(c, v) = \frac{p(c, v)q_V(v)}{\int p(c, v)q_V(v) dv} \propto p(c, v)q_V(v). \quad (76)$$

This is probably the closest aligned to practice, so this is the one we study in this work. To justify this option, imagine the following scenario. If we wish to learn the concept of *red color*, a first step would be to curate a dataset of red objects. To do this, we first consider a collection of photos of objects of varying color and then filter out the ones that look red. The concept conditional measure we define aligns with this process. To learn the actual red concept accurately, our theory predicts that it is sufficient to have additional datasets of objects that are not red, from which we can distinguish red objects, thereby learning the concept of red color.

The next question is how to define the measure  $q_V$ . When considering a single concept  $A^C Z = b^C$  the most natural option to consider  $N(v^C, \sigma^2 \text{Id}_V)$  where  $v^C \in V$  is the unique solution of  $A^C v^C = b^C$  and  $\sigma > 0$  is a positive constant. This is what we do in this work (note that  $\sigma^2$  can be set to 1 by scaling the concept and valuation accordingly).

However, we can also use alternate definitions as suggested above. For instance, we can set  $AZ \stackrel{\mathcal{D}}{=} N(b^C, \text{Id})$ . Then  $Z \sim N(v^C, (A^\top A)^{-1})$ . However, this runs into some technical issues we sketch (and leave to future work to handle this). Consider the intersection of multiple concepts  $C^e$ . In this case the concept space is given by the intersection  $C = \bigcap C^e$  and  $C_0 = \bigcap (C^e)_0$  and we have the orthogonal decomposition  $\mathbb{R}^{d_z} = C_0 \oplus \sum V^e$ . In general the spaces  $V^e$  are not necessarily orthogonal but it is reasonable to assume

that the non-degeneracy condition  $\dim(\sum V^i) = \sum \dim(V^e)$  holds. Now set  $V = \sum V^e$ . If we choose just the standard normal distribution for  $q_{V^e}$  we can define just as in our approach

$$q_V \sim N(v^C, \sigma^2 \text{Id}_V). \quad (77)$$

The second option is to enforce that the marginals of  $q_V$  agree with  $q_{V^e}$ , i.e.,  $q_V(\Pi_{V^e}(v) \in O) = q_{V^e}(O)$  for  $O \subset V^e$ . This results in the set of equations for all  $i$

$$A^e \Sigma (A^e)^\top = \text{Id}_{V^e}. \quad (78)$$

It is likely that this system has a unique solution when non-degeneracy holds for  $V^e$  and this is clearly true for orthogonal spaces but it is not clear how to solve this in general.

## D Inference-Time Intervention of Large Language Models

In this section, we first briefly describe Large Language Models and the recent Inference-Time Intervention (ITI) technique proposed for LLM alignment, which we build on. Then, we use our framework to provide better intuition on some intriguing observations about ITI, including why it works. And then we exploit our ideas to improve the performance of ITI by choosing the steering direction to be a matrix instead of a vector.

### D.1 Preliminaries

**Large Language Models (LLMs)** LLMs are large models capable of generating meaningful text given a context sentence. Due to large-scale training, modern LLMs have shown remarkable capabilities and achieve expert-human-like performance in many benchmarks simultaneously. The architecture of many generative pre-trained transformers (GPT)-style LLMs consists of several transformer layers stacked on top of each other. Since we'll be intervening on them during inference, we'll describe the transformer architecture [102, 21] briefly here. First, the sequence of input tokens (tokens are sub-word units) are encoded into a vector  $x_0$  using a (learned) text embedding matrix and in many cases also a positional embedding matrix. Then, a series of transformer layers act on this vector which passes through a residual stream, to obtain vectors  $x_0, x_1, \dots, x_n$ . The final vector  $x_n$  is then decoded back into token probabilities with a (learned) unembedding matrix. Each transformer layer consists of a multi-head attention mechanism and a standard multilayer perceptron, which captures the nonlinearity.

In the  $l$ th layer, each single multi-head attention mechanism can be described as

$$x_{l+1} = x_l + \sum_{h=1}^H Q_l^h x_l^h, \quad x_l^h = \text{Att}_l^h(P_l^h x_l)$$

Here,  $P_l^h$  and  $Q_l^h$  are matrices that linearly map the vector to an activation space and back respectively, and  $\text{Att}$  denotes the attention mechanism that allows communication across tokens. Here, we have kept the notation consistent with Li et al. [51] for the sake of clarity.

In our setting, we consider the entire set of activations as the learnt latent vector  $Z$ . That is, the input is  $x = x_0$  and the pre-trained model is essentially the function  $f$  such that  $f(x)$  consists of the concatenation of the vectors  $\{x_l\}_{l \geq 1}$ , the intermediate activations  $\{x_l^h\}_{l \geq 0}$  and also the output of the linear transformations  $\{P_l^h x_l\}_{l \geq 0}, \{Q_l^h x_l^h\}_{l \geq 0}$ . Our theory hinges on the assumption that pre-trained LLMs satisfy the linear representation hypothesis, that is, various relevant concepts can be realized via linear transformations of the latent transformation  $f(x)$ . Indeed, this has been empirically observed to hold in many prior works [12, 96, 65, 63, 51, 68, 30, 40] (see also related works on geometry of representations [39, 40] and references therein). It's a fascinating question why such models trained with next token prediction loss also learn linear representations of various human-interpretable concepts such as sentiment, see Jiang et al. [40] for recent progress on this problem.

It's well-known that despite large-scale pretraining and subsequent improvement of pre-trained models via techniques like Reinforcement Learning with Human Feedback (RLHF) and Supervised Fine-Tuning (SFT) [67, 5, 97], significant issues still remain [86], e.g., the model can hallucinate or generate incorrect responses (even though the model *knows* the correct response which can be extracted via other means, e.g., Chain-of-Thought prompting [108]). Various methods have been proposed to fine-tune the models [67, 5, 6, 97, 75] but many of them are expensive and time- and resource-intensive as they requires huge annotation and computation resources. Therefore, more efficient techniques are highly desired, one of which is the category of methods known as activation patching. activation patching (also called activation editing or activation engineering) [31, 105, 91, 99, 116, 111, 50, 60].

**Inference-Time Intervention, an activation patching method for truthfulness** Activation patching is a simple minimally invasive technique to align LLMs to human-preferences. Specifically, given various concepts such as truthfulness, activation patching makes modifications to the model during inference time so that the desired concepts can be aligned. This technique can be thought of as an application of the emerging field of mechanistic interpretability [66], which aims to interpret the learnt latent vector in terms of human-interpretable concepts, thereby allowing us to reverse-engineer what large models learn.

Activation patching has many variants [50, 31, 60], but we’ll focus on the simple technique of adding *steering vectors* to various intermediate layers during intervention [91, 99, 51, 80]. This means that during inference, the output activations are modified by adding a constant vector in order to promote alignment of some concept. The vector will be learnt independently based on separate training data.

In particular, a recent technique called Inference-Time Intervention (ITI) was proposed to do this for the specific concept of truthfulness. ITI focuses on the activation heads  $\{\text{Att}_l^h(P_l^h x_l)\}_{l \geq 0}$  and add to them steering vectors in order to promote truthfulness. To learn the steering vectors, a subset of the TruthfulQA dataset [53], namely a dataset of questions  $q_i$  with annotated true ( $a_{i,j}, 0$ ) and false answers ( $a_{i,j}, 1$ ), are prepared as  $\{q_i, a_i, y_i\}_{i=1,2,\dots}$ . For each sample, the question and answer are concatenated as a pair and the corresponding activations of the heads  $x_l^h$  (for the final token) are computed via forward passes. Then, a linear probe sigmoid( $\langle \theta, x_l^h \rangle$ ) is independently trained on each activation head to distinguish true from false answers. Finally, the top  $K$  heads based on the accuracy of this classification task are chosen (for a tunable hyperparameter  $K$ ) and the steering vector  $\theta_l^h$  for the  $h$ -th head in layer  $l$  is chosen to be the mean difference of the activations between the true and false inputs. The intuition is that this direction roughly captures the direction towards truthfulness.

Formally, for the  $h$ th head of the  $l$ th layer, ITI adds the steering vector  $\alpha \sigma_l^h \theta_l^h$  so as to get

$$x_{l+1} = x_l + \sum_{h=1}^H Q_l^h(x_l^h + \alpha \sigma_l^h \theta_l^h), \quad x_l^h = \text{Att}_l^h(P_l^h x_l)$$

during inference. Here,  $\theta_l^h$  is the steering vector,  $\sigma_l^h$  is the standard deviation of the activations of this head along the chosen direction and  $\alpha$  is a hyperparameter. That is, the activations are shifted along the truthful directions by a multiple of the standard deviation, and this is repeated autoregressively. Note that this does not depend on the specific GPT-like model being used. The intuition is that during inference, the activations are intervened upon to shift towards the truthful direction. The top  $K$  heads are chosen to be minimally intrusive and also a design choice based on observations of the probing metrics.

**Performance of ITI** In Li et al. [51], ITI was shown to significantly improve the truthfulness of various LLMs after having been trained on as few as a few dozen samples, compared to what’s needed for Reinforcement Learning based techniques [67, 26]. ITI was evaluated on the TruthfulQA benchmark [53], which is a hard adversarial benchmark to evaluate truthfulness of language models. In particular, it contains 817 questions with a multiple-choice and generation tracks, spanning 38 categories such as logical falsehoods, conspiracies and common points of confusion. For the multiple-choice questions, the accuracy is determined by the conditional probabilities of candidate answers given the question. Evaluating the generation track questions is harder, and it is done by generating a model output and then evaluating it via a finetuned GPT-3-13B model [53, 64]. Moreover, the choice of the intervention strength  $\alpha$  is calibrated so that it’s neither too small (to promote truthfulness) nor too large (to ensure the original capabilities of the LLM are not lost). To check if the original capabilities are preserved, [51] compute two additional quantities to measure how far the modified model deviates from the original model. These are the Cross-Entropy (CE) loss, which is standard in language modeling and the Kullback–Leibler divergence (KL div.) of the next token probabilities before and after intervention. To compute these quantities, a subset of Open Web Text is used [73]. Finally, it was shown that ITI implemented on the LLaMA [97], Alpaca [94] and Vicuna [14] models significantly improved their performance on the TruthfulQA benchmark compared to the baseline models. Moreover, in many cases, it also beat other techniques such as few-shot prompting and supervised fine-tuning. Please see Li et al. [51] for additional details.

## D.2 Interesting observations of ITI

While the elegant ITI technique was designed to align LLMs towards truthfulness in practice, it also raised fascinating and intriguing questions in mechanistic interpretability. In addition to improving the technique of ITI itself, our work makes progress towards some of these questions via our framework.

1. The authors of Li et al. [51] state in section 2 that although the technique works well in practice, it's not clear what ITI does to the model's internal representations. In addition, prior works [12, 96, 65, 63, 68, 40] have observed empirically that the latent representations learned by LLMs seem to have interpretable linear directions, which ITI exploits. We use our framework to illustrate in more detail one possible explanation of what ITI does to the model representations and why it works, in the next section.
2. The authors visualize the geometry of “truth” representations in section 3.2 of their work via the following experiment: For the most significant head (layer 14, head 18), after finding the first truthful direction via the linear probing technique, they remove it and attempt to find a second probe orthogonal to the first. They find surprisingly that the second probe is also very informative, leading them to predict that the concept of “truth” lies in a subspace, not a single direction. Restated in our framework, the concept of truthfulness is a non-atomic concept (as per Definition 2). This served as an inspiration for our proposed technique in the next section, where we propose to use steering matrices instead of steering vectors for LLM alignment.
3. As  $\alpha$  was increased, the authors observed that truthfulness of the model increased however helpfulness decreased. This suggests that the “truthfulness” and “helpfulness” concepts are not atomic (as per Definition 2) however they share certain atomic concepts. We leave to future work the exciting question of mechanistically extracting such common atomic concepts.

### D.3 The choice of the steering vector

We will now analyze the truthfulness concept via our framework and give more insight on why the mean of the differences is a reasonable choice of steering vector for ITI. Based on our theory, we will then provide a modification to this choice that uses steering matrices instead of steering vectors. Since this section is based on heuristics and informal assumptions, we will refrain from making any formal claims or analyses. Indeed, a formal analysis of concepts in natural language is a hard problem in general and we do not attempt it here. We conclude with ideas for potential extensions that're worth exploring in future work.

Denote the function  $h$  to be the sequence of head activations  $h(x) = (x_l^h)_{l,h} \in \mathbb{R}^d$ . Note that while we can study general steering vectors for the entire latent space of representations  $f(x)$  learned by LLMs as some works do, ITI focuses only on steering the head activations  $h(x)$ , so we will apply our framework to this subset representation space. In addition, we will make the simplification that we neglect the effects of the steering vector from bottom layers towards the top layers, which we do because we are dealing with sparse steering vectors and also, each single head shift is minor and does not in isolation change the behavior of the model as verified by experiments [51][Appendix B.1].

Applying our framework, we model the concept of truth via the concept matrix  $A \in \mathbb{R}^{d_C \times d}$  and two valuations  $b_0, b_1 \in \mathbb{R}^{d_C}$  corresponding to *False* and *True* respectively. In other words, the set of false sentences and true sentences lie respectively in

$$\mathcal{S}_{\text{false}} = \{x | Ah(x) = b_0\}, \quad \mathcal{S}_{\text{true}} = \{x | Ah(x) = b_1\}$$

Note that they only approximately lie in these spaces because of our notion of concept conditional distribution. However, if we reasonably assume that the Gaussian concentration region is much smaller than the separation between these hyperplanes, then the rest of the arguments in this section should apply.

Now, a steering vector  $\eta$  is a vector such that it moves the activations from the false space to the true space, while keeping other concepts unaffected. That is, if we pick a false sentence  $x$ , i.e.,  $Ah(x) = b_0$ , then the steering vector  $\eta \in \mathbb{R}^d$  essentially steers the activations so that  $A(h(x) + \eta) = b_1$ . In other words, it moves the sentence from false to true. Indeed, many vectors  $\eta$  do satisfy this equality, therefore the goal is to find an optimal  $\eta$  that does not (significantly) affect other concepts of interest. Indeed, the ideal steering vector will be  $A^+(b_1 - b_0)$  where  $A^+$  is the pseudoinverse of  $A$ . This vector will precisely affect this concept space and will not affect the concept valuations for any concept orthogonal to  $A$ . However, the issue is that we do not know  $A$  and therefore we will approximate this steering vector from training samples.

To this end, suppose we are given a collection of counterfactual sentence pairs  $c_i^F, c_i^T$  which correspond to a false answer and a true answer for the same question  $q_i$ . Consider the  $i$ th counterfactual pair  $c_i^F, c_i^T$ . We will assume the reasonable scenario that the only difference among their concepts is the concept of truthfulness. That is, for any other concept  $B$  that is orthogonal to  $A$ , the valuations of these pairs are

identical. Suppose we stack all other (orthogonal) concepts of relevance for this particular counterfactual sample into a matrix  $B_i$ , then we have

$$Ah(c_i^F) = b_0, Ah(c_i^T) = b_1, \quad B_i h(c_i^F) = B_i h(c_i^T)$$

for all  $i = 1, 2, \dots$ .

Now, it's clear that for a single sample  $i$ , we cannot choose  $\eta_i = h(c_i^T) - h(c_i^F)$  as a steering vector. This is because for such a choice of steering vector, while it's still the case that  $B_i(h(c_i^F) + \eta_i) = B_i(h(c_i^T) + \eta_i)$ , it need not be the case that  $B_j(h(c_j^F) + \eta_i) = B_j(h(c_j^T) + \eta_i)$  for some  $j \neq i$ , i.e., not all orthogonal concepts are preserved. This happens because  $\eta_i = h(c_i^T) - h(c_i^F)$  not only contains the correct shift along the span of  $A$  but also unnecessarily shifts along the space orthogonal to  $B_i$ , which will affect  $B_j$ .

However, when we have a large enough number of contexts  $j$  with sufficient concept diversity, we can make the reasonable prediction that since the context questions  $q_j$  are sufficiently diverse (therefore  $B_j$  are diverse), the corresponding projections  $B_j \eta_i$  are essentially uncorrelated vectors as  $i$  varies. In this case, we can pick  $\eta$  to be the mean of all the shift vectors  $\eta_i$  to achieve the desired effect. Therefore, under the above simplifying assumptions, choosing  $\eta$  to be the mean of  $\eta_1, \eta_2, \dots$  will satisfy

$$\begin{aligned} Ah(x) &= b_0, A(h(x) + \eta) = b_1 && \text{exact equalities for truth concept } A \\ Bh(x) &\approx B(h(x) + \eta) && \text{approximate equalities for concepts } B \text{ orthogonal to } A \end{aligned}$$

for any concept  $B$  orthogonal to  $A$ . This explains why the choice of mean of the activation differences across counterfactual pairs is a reasonable choice of steering vector. This is precisely the technique used in ITI. While they also experiment with other steering vectors, they found that this works the best for their experiments.

Now, we will continue on our insights to analyze whether we can build better steering vectors  $\eta$ . We present two crucial insights based on our analysis so far.

1. Looking at our desired equations, any *weighted combination* of  $\eta_i = h(c_i^T) - h(c_i^F)$  will satisfy  $Ah(x) = b_0, A(h(x) + \eta) = b_1$  exactly.
2. We could potentially choose the steering vector  $\eta$  to be a function of  $x$  instead of being a constant vector, provided  $\eta(x)$  is efficiently computable during inference time.

Exploiting our first insight, we conclude that choosing any weighted combination of the  $\eta_i$  should be a reasonable choice of steering vector provided we can control its effects on the spaces orthogonal to  $A$ . That is, we can choose

$$\eta = \sum_i w_i \eta_i = \sum_i w_i (h(c_i^T) - h(c_i^F))$$

as our steering vector. This gives us the extra freedom to tune the weights  $w_1, w_2, \dots$  based on other heuristics. Note that this also captures the choice of the top principal component of the steering vector as experimented in [96].

Our second observation suggests that even the steering vector  $\eta$  could be a function of  $x$ , namely  $\eta(x)$ , provided it's efficiently computable during inference. Therefore, this suggests the usage of

$$\eta(x) = \sum_i w_i(x) (h(c_i^T) - h(c_i^F))$$

as our steering vector where the weights  $w_i(x)$  depend on  $x$ .

Based on these two observations, we propose our ITI modification. We choose the steering vector to be dependent on the context  $x$ , with weights chosen to be  $w_i = \langle \lambda(x), \lambda(c_i^F) \rangle$  for a sentence embedding  $\lambda$  (such as Sentence-BERT [79]). That is,

$$\eta(x) = \sum_i \langle \lambda(x), \lambda(c_i^F) \rangle (h(c_i^T) - h(c_i^F))$$

Indeed, this is reasonable as if a context  $x$  is close to  $c_i^F$  for a specific training sample  $i$  in terms of their sentence embeddings  $\lambda(x)$  and  $\lambda(c_i^F)$ , then this particular sample's steering vector should be upsampled. In other words, we can think of the training sample contexts as voting on their respective counterfactual steering vector, with weights determined by the similarity between the representation

of the test context and the representation of the sample context. A justification would be that  $B_i(x)$  (the relevant concepts for a datapoint) depend smoothly on  $x$  (proximity is measured by similarity of embeddings) so it makes sense to upweight close points.

Finally, we need to argue that we can compute this efficiently during inference. For this, we exploit the structure of our steering vector representation as follows.

$$\begin{aligned}\eta(x) &= \sum_i \langle \lambda(x), \lambda(c_i^F) \rangle (h(c_i^T) - h(c_i^F)) \\ &= \left( \sum_i (h(c_i^T) - h(c_i^F)) \lambda(c_i^F)' \right) h(x) \\ &= Mh(x)\end{aligned}$$

for the matrix  $M = \sum_i (h(c_i^T) - h(c_i^F)) \lambda(c_i^F)'$ , where  $v'$  denotes the tranposed vector. We remark that the weights  $w_i(x)$  as used could potentially be negative but this is not an issue since the projection of the corresponding counterfactual vector in the direction of  $B$  is still random and we finally normalize  $\eta(x)$ , so the magnitude doesn't matter.

Therefore, this steering can be done efficiently by precomputing the *steering matrix*  $M$  and then during inference, we simply compute the steering vector  $\eta(x)$  as  $\eta(x) = Mh(x)$ .

**Implementation considerations** We briefly note down some design choices we made in our implementation of the above method.

1. Since  $\eta(x)$  is a function of  $x$ , the standard deviation of the activation projection on this direction, i.e.,  $\sigma_l^h(x)$  cannot be precomputed (as Li et al. [51] do), therefore we compute them dynamically during inference, which takes little overhead with fast tensorization operations (in particular, this is not the slow step).
2. We opted to go with evaluating the model only on the multiple-choice questions. This is partly because to evaluate the generated text, the recommended method is to use fine-tuned GPT-3-13B models but OpenAI have retired many of their older models as of this year, and therefore, the entire batch of experiments would have to be rerun with their newer models which could potentially change the baselines, and also because this work is a proof-of-concept rather than a comprehensive evaluation.
3. For computing the sentence embeddings, we only use the question prompts, as they contain all relevant contexts. And we normalize  $\eta(x)$  during inference time.

**Additional ideas for improvement** We re-iterate that our experimental exploration is not exhaustive and the preliminary experiments are merely meant to be a proof-of-concept. In this section, building on our insights, we outline some further ideas to improve the performance of ITI. We leave to future work to comprehensively explore these techniques in order to extract better performance towards LLM alignment.

1. Note that we opted to go with the weights  $\langle \lambda(x), \lambda(c_i^F) \rangle$  where  $\lambda$  was chosen to be a sentence transformer embedding [79]. While this is a reasonable choice, similarity metrics could be measured in other ways, e.g., with other sentence embedding models.
2. Going further, the weights do not have to be similarity scores and could be chosen via other heuristics. For instance, they could be chosen to be constants but potentially be optimized using a hold-out test set.
3. As Li et al. [51] noted, the ITI technique could be applied on top of fine-tuned models in order to further improve their performance. Therefore, our proposed modification could also potentially be applied on top of fine-tuned models.

## E Contrastive algorithm for end-to-end concept learning

In this section, we will describe the details of the contrastive learning method for learning concept representations end-to-end. First, we will prove the computation of the true log-odds.

**Lemma 2.** For any concept index  $e$ , there exist some constants  $c_e$  such that

$$\begin{aligned} & \ln(p^e(Z)) - \ln(p(Z)) \\ &= \sum_{i=1}^n \left( -\frac{1}{2} M_{ei} \langle a_i, Z^e \rangle^2 + B_{ei} \langle a_i, Z^e \rangle \right) + c_e \end{aligned}$$

where  $M, B$  are the environment-concept matrix and the environment-valuation matrix defined in (7) and (8).

*Proof.* This follows from Eq. (13) in the proof of Theorem 2.  $\square$

From our main identifiability results, we can assume without loss of generality that the concept vectors we learn are coordinate vectors. In other words, we consider a neural network  $h^\theta$  with parameters  $\theta$  with output neurons  $h_1^\theta, \dots, h_n^\theta$  such that the  $n$  atomic concepts will now correspond to the concept vectors  $e_1, \dots, e_n$  (which is reasonable as they are only identifiable up to linear transformations). Therefore, for each environment  $e$ , we can train classifiers of the form

$$g_e(X, \alpha^e, \beta_k^e, \gamma_k^e, \theta) = \alpha^e - \sum_{k=1}^{\dim(C_e)} (\beta_k^e h_k^\theta(X))^2 + \sum_{k=1}^{\dim(C_e)} \gamma_k^e (h_k^\theta(X))$$

equipped with standard cross-entropy loss, for hyperparameters  $\alpha^e, \beta_k^e, \gamma_k^e, \theta$ . Indeed, this is reasonable since if the training reaches the global optima in the ideal case, then the loss function will correspond to the Bayes optimal classifier and therefore,  $g_e(X, \alpha^e, \beta_k^e, \gamma_k^e, \theta) = \ln(p^e(Z)) - \ln(p(Z))$ , which along with Lemma 2 will suggest that the learnt network  $h$  is linearly related to the function  $A^e f^{-1}$ , as desired. Lastly, we choose the loss function to be the aggregated CE loss and an extra regularization term. That is,

$$\mathcal{L} = \sum_e \underbrace{-\mathbb{E}_{j \sim \text{Unif}(\{0, e\})} \mathbb{E}_{X \sim X^e} \left( \ln \frac{e^{\mathbf{1}_{j=e} g_e(X)}}{1 + e^{g_e(X)}} \right)}_{\text{CE loss for environment } e} + \eta \|\beta\|_1$$

for a regularization hyperparameter  $\eta$ .

## F Additional details about the synthetic setup

In this section, we detail our synthetic setup in Section 5.1. The base distribution is sampled from a Gaussian mixture model with 3 components whose parameters are chosen randomly. The weights are randomly chosen from  $\text{Unif}(0.3, 1)$  (and then normalized), the entries of the means are chosen from  $\text{Unif}(-1, 1)$  and the covariance is chosen to be a diagonal matrix with entries in  $\text{Unif}(0.01, 0.015)$  (note that the diagonal nature doesn't really matter since a map  $f$  will be applied to this distribution). The mixing function  $f$  is chosen to be either (i) linear or (ii) nonlinear with a 1-layer MLP containing 16 hidden neurons and LeakyReLU(0.2) activations.

The number of concepts  $n$  is intentionally chosen to be less than the ground truth dimension  $d_z$  and the number of concepts is  $m = n + 1$  as per our theory. The concepts are taken to be atomic, with the concept vectors and valuations chosen randomly, where each entry of the concept vector is chosen i.i.d from  $\text{Unif}(-0.3, 0.3)$ , and the resampling distribution is chosen to be a Gaussian with variance 0.005. Finally, we choose 5000 samples per environment, sampled via the rejection sampling Algorithm 1. For the contrastive algorithm, we choose the architecture to either be linear or nonlinear with a 2-layer MLP with 32 hidden neurons in each layer, with the final parametric layer chosen based on the known concept, to have the form described above. We train for 100 epochs with  $\eta = 0.0001$  and use Adam optimizer with learning rates 0.5 for the parametric layer and 0.005 for the non-parametric layer, with a Cosine Annealing schedule [57].

## G Controllable generative modeling via rejection sampling

In this section, we will describe how to sample from a concept conditional distribution with a known concept. Once the concepts are learned in our framework, we can use this technique to generate new data satisfying various desired concepts, which will aid in controllable generative modeling.

Consider the base distribution on  $Z \in \mathbb{R}^{d_z}$  with density  $p(Z)$ . Suppose we wish to sample from a concept  $C$  given by  $AZ = b$  and resampling distribution  $q$ . We additionally assume that  $q$  is efficiently computable and an upper bound  $L$  is known for its density, i.e.,  $L \geq \max(q)$ .

Recall that the desired density is defined as

$$p_C(Z) \propto p(Z) \prod_{i \leq \dim(C)} q((AZ - b)_i)$$

Note that it's infeasible to compute the normalization constant for such complex distributions. However, we bypass this by using rejection sampling. We describe the procedure in Algorithm 1.

---

**Algorithm 1:** Rejection sampling for controllable generative modeling

---

**Input:**

- Base distribution  $p$
- Resampling distribution  $q$  with upper bound  $L \geq \max(q)$
- Concept  $C$  with transformation  $A$  and valuation  $C$

**Output:** Returns a single sample from  $p_C(Z)$

```

1  $M = L^{\dim(C)}$ 
// Repeat trials until condition is met
2 while True do
3   Z = yield(p)
4   U = yield(Unif(0, 1))
5    $R = \frac{1}{M} \prod_{i \leq \dim(C)} q((AZ - b)_i)$ 
6   if  $R \geq U$  then
7     return Z

```

---

Informally, we first sample  $Z \sim p$  (we overload notation for both density and the distribution) and an independent variable  $U \sim Unif(0, 1)$ , the uniform distribution on  $(0, 1)$ . We accept the variable  $Z$  if

$$\frac{1}{M} \prod_{i \leq \dim(C)} q((AZ - b)_i) \geq U$$

for a predetermined upper bound  $M$  on the quantity  $\prod_{i \leq \dim(C)} q((AZ - b)_i)$ . If the inequality is false, we simply reject the sample and repeat.

Now we will argue why this algorithm is correct, which is accomplished in Theorem 3. Let

$$N_C = \int_Z p(Z) \prod_{i \leq \dim(C)} q((AZ - b)_i)$$

be the normalization constant in the definition of  $p_C(Z)$ . Therefore

$$p_C(Z) = \frac{1}{N_C} p(Z) \prod_{i \leq \dim(C)} q((AZ - b)_i)$$

**Lemma 4.** Let  $M \geq \max(q)^{\dim(C)}$ . The acceptance probability of each iteration of the while loop in Algorithm 1 is  $\Pr[Z \text{ accepted}] = \frac{N_C}{M}$

*Proof.* We have

$$\begin{aligned} \Pr[Z \text{ accepted}] &= \Pr_{U,Z} \left[ U \leq \frac{1}{M} \prod_{i \leq \dim(C)} q((AZ - b)_i) \right] \\ &= \Pr_{U,Z} \left[ U \leq \prod_{i \leq \dim(C)} \frac{q((AZ - b)_i)}{\max(q)} \right] \quad \text{since } M \geq \max(q)^{\dim(C)} \end{aligned}$$

$$\begin{aligned}
&= \int_Z \Pr_U \left[ U \leq \prod_{i \leq \dim(C)} \frac{q((AZ - b)_i)}{\max(q)} \right] p(Z) dZ \quad \text{as } U, Z \text{ are independent} \\
&= \int_Z \left[ \prod_{i \leq \dim(C)} \frac{q((AZ - b)_i)}{\max(q)} \right] p(Z) dZ \quad \text{since } \frac{q((AZ - b)_i)}{\max(q)} \leq 1 \text{ always} \\
&= \int_Z \frac{N_C p_C(Z)}{M} dZ \\
&= \frac{N_C}{M}
\end{aligned}$$

□

Before we prove correctness, we will remark on the expected number of trials needed for accepting each sample.

**Corollary 1.** *The expected number of trials needed to generate a single sample is  $\frac{M}{N_C}$*

*Proof.* Note that each iteration of the while loop is independent, therefore the number of trials until acceptance is distributed as a geometric random variable whose expectation is the inverse of the parameter.

□

This suggests that for our algorithm to be efficient in practice,  $M$  should be chosen as small as possible, i.e., estimates of  $\max(q)$  should be as tight as possible.

**Theorem 3.** *Algorithm 1 yields samples from the concept conditional distribution  $p_C$ .*

*Proof.* The proof is at heart the proof of correctness of rejection sampling. For arbitrary parameters  $t_1, \dots, t_{d_z} \in \mathbb{R}$ , let's compute the cumulative density of the samples output by Algorithm 1 and show that it matches the cumulative distribution function of  $p_C(Z)$  evaluated at  $t_1, \dots, t_{d_z}$ , which will complete the proof. That is, we wish to calculate

$$\Pr[Z_1 \leq t_1, \dots, Z_{d_z} \leq t_{d_z} | Z \text{ accepted}] = \frac{\Pr[Z_1 \leq t_1, \dots, Z_{d_z} \leq t_{d_z}, Z \text{ accepted}]}{\Pr[Z \text{ accepted}]}$$

We already computed the denominator in Lemma 4. Therefore,

$$\begin{aligned}
&\Pr[Z_1 \leq t_1, \dots, Z_{d_z} \leq t_{d_z} | Z \text{ accepted}] \\
&= \frac{M}{N_C} \Pr[Z_1 \leq t_1, \dots, Z_{d_z} \leq t_{d_z}, Z \text{ accepted}] \\
&= \frac{M}{N_C} \mathbb{E}_Z [\mathbb{1}_{Z_1 \leq t_1} \dots \mathbb{1}_{Z_{d_z} \leq t_{d_z}} \cdot \mathbb{E}_U [\mathbb{1}_{Z \text{ accepted}}]] \\
&= \frac{M}{N_C} \mathbb{E}_Z \left[ \mathbb{1}_{Z_1 \leq t_1} \dots \mathbb{1}_{Z_{d_z} \leq t_{d_z}} \cdot \frac{1}{M} \prod_{i \leq \dim(C)} q((AZ - b)_i) \right] \quad \text{from the proof of Lemma 4} \\
&= \int_Z \mathbb{1}_{Z_1 \leq t_1} \dots \mathbb{1}_{Z_{d_z} \leq t_{d_z}} \cdot \frac{1}{N_C} \prod_{i \leq \dim(C)} q((AZ - b)_i) p(Z) dZ \\
&= \int_Z \mathbb{1}_{Z_1 \leq t_1} \dots \mathbb{1}_{Z_{d_z} \leq t_{d_z}} \cdot p_C(Z) dZ
\end{aligned}$$

which is precisely the cumulative distribution function of  $p_C(Z)$  evaluated at  $t_1, \dots, t_{d_z}$ . □