

# Causal Imputation for Counterfactual SCMs: Bridging Graphs and Latent Factor Models

**Álvaro Ribot\***

ARIBOTBARRADO@G.HARVARD.EDU

*School of Engineering and Applied Sciences, Harvard University, and CFIS, UPC*

**Chandler Squires**

CSQUIRES@MIT.EDU

*Laboratory for Information and Decision Systems, MIT, and Broad Institute of MIT and Harvard*

**Caroline Uhler**

CUHLER@MIT.EDU

*Laboratory for Information and Decision Systems, MIT, and Broad Institute of MIT and Harvard*

**Editors:** Francesco Locatello and Vanessa Didelez

## Abstract

We consider the task of *causal imputation*, where we aim to predict the outcomes of some set of actions across a wide range of possible contexts. As a running example, we consider predicting how different drugs affect cells from different cell types. We study the *index-only* setting, where the actions and contexts are categorical variables with a finite number of possible values. Even in this simple setting, a practical challenge arises, since often only a small subset of possible action-context pairs have been studied. Thus, models must extrapolate to novel action-context pairs, which can be framed as a form of matrix completion with rows indexed by actions, columns indexed by contexts, and matrix entries corresponding to outcomes. We introduce a novel SCM-based model class, where the outcome is expressed as a counterfactual, actions are expressed as interventions on an instrumental variable, and contexts are defined based on the initial state of the system. We show that, under a linearity assumption, this setup induces a *latent factor model* over the matrix of outcomes, with an additional fixed effect term. To perform causal prediction based on this model class, we introduce simple extension to the Synthetic Interventions estimator (Agarwal et al., 2020). We evaluate several matrix completion approaches on the PRISM drug repurposing dataset, showing that our method outperforms all other considered matrix completion approaches.

**Keywords:** Causal imputation, latent factor models, synthetic interventions, matrix completion

## 1. Introduction

A core goal in scientific modeling – often left implicit – is to construct models that accurately predict a system’s behavior across a wide range of conditions. More precisely, we consider the following general *causal prediction* problem:

*We know some context  $C$ , i.e., some partial information about the system’s state.*

*We are considering whether to perform an action  $A$  that will affect the system.*

*We wish to predict some outcome  $Y$ , i.e. some feature(s) of the system after performing  $A$ .*

---

\* Á. Ribot’s contributions to this work were made while visiting MIT IDSS and while affiliated with Centre de Formació Interdisciplinària Superior (CFIS) - Universitat Politècnica de Catalunya (UPC).

$$\mathbf{Y} = \left. \begin{pmatrix} Y_{11} & Y_{12} & ? & \cdots & Y_{1n} \\ ? & ? & Y_{23} & \cdots & Y_{2n} \\ Y_{31} & ? & Y_{33} & \cdots & ? \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ Y_{m1} & ? & ? & \cdots & ? \end{pmatrix} \right\} \begin{array}{l} \text{actions } i \in [m] \\ \text{contexts } j \in [n] \end{array}$$

Figure 1: **The causal matrix completion problem.** Each row of  $\mathbf{Y}$  corresponds to an action, and each column corresponds to a context.  $Y_{ij}$  denotes the outcome after performing action  $i$  in context  $j$ . We represent missing entries with "?".

In general, causal prediction may involve high-dimensional contexts, actions, and/or outcomes, such as images or text (Chalupka et al., 2015; Castro et al., 2020; Feder et al., 2022). In this work, we study causal prediction in the *index-only* setting<sup>1</sup>, where actions take values in  $[m] := \{1, \dots, m\}$  and contexts take values in  $[n] := \{1, \dots, n\}$ , with the values  $i \in [m]$  and  $j \in [n]$  carrying no semantic meaning. In particular, the only information available about an action/context is an index which distinguishes it from other actions/contexts; there is no prior notion of similarity between actions/contexts. To emphasize our setting, we will write  $I_A$  instead of  $A$  and  $I_C$  instead of  $C$ . By restricting our focus to the index-only setting, we aim to maximize the clarity of this work, while building a solid foundation for future works.

As a running example of causal prediction, we consider the problem of *viability prediction* in biology, an essential step in tasks such as drug repurposing (Wu et al., 2022). In this problem, we are given a cell type  $j \in [n]$  (e.g., skin or lung) and a drug label  $i \in [m]$  (e.g. tyloxapol or gefepirine), and we aim to predict what proportion  $Y$  of cells will survive if we administer drug  $i$  to a large group of type- $j$  cells. We note that existing databases (Southan et al., 2013) offer extensive additional information about drugs, which can be used when extending beyond the index-only setting considered here. Thus, the prediction accuracy in our setting should be seen as a lower bound on the accuracy that could be attained by leveraging this additional information.

To generate causal predictions, we need some model specifying how an action  $I_A$  interacts with the context  $I_C$  to produce the outcome  $Y$ . In statistics and machine learning, a model is obtained via *data-driven* approaches, which consider a fixed *model class*  $\Theta$  and use data to select a model  $\hat{\theta}$  from  $\Theta$  (called *learning* or *model fitting*). In this work, we consider *supervised* learning, where the available data consists of samples of  $(I_A, I_C, Y)$ . For large  $m$  and  $n$ , the available data often contains samples from only a small subset  $\Omega$  of all possible  $m \cdot n$  possible action-context pairs. For example, given  $n \approx 100$  cell types and  $m \approx 10,000$  drugs, we have  $m \cdot n \approx 1$  million, but our dataset might cover only 5-10% of these pairs. By arranging the available data into a matrix with rows indexed by actions and columns indexed by contexts, we obtain a partially-observed matrix  $\mathbf{Y}$ ,

1. This setting is similar to the *tabular* setting in reinforcement learning; we choose to use the term *index-only* to emphasize the lack of any additional structure.




Figure 2: **The latent factor model (LFM) written as a simple structural causal model (SCM).** LFMs assume that  $Y_{ij} = \langle \mathbf{u}_i, \mathbf{v}_j \rangle + \varepsilon_{ij}$  for  $\mathbf{u}_i, \mathbf{v}_j \in \mathbb{R}^r$ . This generative process can be viewed as an SCM over 3 observed (shaded) and 2 latent (unshaded) variables. Here,  $I_A$  and  $I_C$  represent action and context indices, respectively (note that they are independent).

as in Fig. 1. In these situations, causal prediction requires *extrapolating* from  $\Omega$  to the entire space  $[m] \times [n]$ , a problem known as *causal imputation* (Squires et al., 2022).

For extrapolation to be feasible, the model class  $\Theta$  must encode some inductive biases governing the interplay between actions, contexts, and outcomes. The model class must also be reasonably well-specified, i.e., there must exist some  $\theta^* \in \Theta$  which accurately describes the relationship between  $Y$ ,  $I_A$ , and  $I_C$ , at least to a good approximation. Given the importance of the model class  $\Theta$ , this work aims to advance our understanding of the relationship between two model classes that are commonly used for causal prediction. In particular, we study the class of *latent factors models* (LFMs) and the class of *structural causal models* (SCMs).

LFMs, also known as *interactive fixed effects models*, are widely used in econometrics and recommendation systems (Athey et al., 2021; Koren et al., 2009). In an LFM, each action  $i \in [m]$  is associated with an unknown vector  $\mathbf{u}_i \in \mathbb{R}^r$ , and each context  $j \in [n]$  is associated with an unknown vector  $\mathbf{v}_j \in \mathbb{R}^r$ . LFMs assume  $Y_{ij} = \langle \mathbf{u}_i, \mathbf{v}_j \rangle + \varepsilon_{ij}$  for some mean-zero  $\varepsilon_{ij}$ , and can be seen as a simple form of SCM, see Fig. 2. We can also extend these models to include terms for a *fixed action effect* and/or a *fixed context effect* by letting  $Y_{ij} = \langle \mathbf{v}_i, \mathbf{u}_j \rangle + \alpha_i + \beta_j + \varepsilon_{ij}$  for some  $\alpha_i, \beta_j \in \mathbb{R}$ . LFMs are related to low-rank factorizations: letting  $\mathbf{L} \in \mathbb{R}^{m \times n}$  with  $\mathbf{L}_{ij} = \mathbb{E}[Y_{ij}]$ , the definition of the LFM (with no fixed effects) implies that  $\text{rank}(\mathbf{L}) \leq r$ .

LFMs are often assumed as a model class without any “deeper” justification. One intuitive way to motivate LFMs is to show that they arise from other (potentially more transparent) modeling assumptions, as in Udell and Townsend (2017). Relating LFMs to other model classes is important for several reasons: such connections offer insights that could legitimize trust in the model’s prediction, and can serve as a starting point from which to develop more general model classes. Thus, a primary aim of this work is to offer a new justification for LFMs, starting from the assumption that the system’s state obeys a structural causal model (SCM), with  $Y_{ij}$  defined as a *counterfactual* outcome under action  $i$  when the system is in context  $j$ . That is, in this paper, we work with quantities that are defined in terms of a counterfactual distribution, and which cannot be defined only using interventional distributions. Such quantities are commonly seen, for example, in the literature on mediation analysis (Malinsky et al., 2019).

Finally, note that there is a substantial difference between the roles that columns (contexts) and rows (actions) are playing. Inspired by this distinction, let us call a matrix completion approach *symmetric* when, performed on  $\mathbf{Y}^\top$ , we obtain the same predictions. We want to use a non-symmetric approach that is compatible with our causal viewpoint.

**Organization of the paper and contributions.** After an overview of related work in Section 2, we review SCMs and formally define our model class in Section 3. We specialize to linear models in Section 4, where we establish the following results:

- We generalize the model class from Squires et al. (2022) to allow for the context  $I_C$  to be defined in terms of the system’s state, which was not previously allowed. In Theorem 1, we show that our new SCM-based model class implies an LFM with fixed action effects.
- We propose a slight modification to the *Synthetic Interventions* estimator of Agarwal et al. (2021) that accounts for fixed action effects.
- We give additional conditions on contexts and actions, which imply further structure on  $\mathbf{L}$ ; see Prop. 2 and Corollary 3. We show how this additional structure can be leveraged to improve estimation and hypothesis testing.

Finally, in Section 5, we compare the performance of several causal prediction approaches on the PRISM drug repurposing dataset, showing that our method outperforms alternative approaches.

## 2. Related Work

Since this paper bridges between latent factor models and structural causal models, the potential scope for related work is vast. We limit our discussion to works that consider the index-only setting. These works are complementary to recent works which give extrapolation guarantees for more general/structured action spaces, such as Saengkyongam et al. (2023) and Agarwal et al. (2023). A summary of the algorithms presented can be found in Table 1. We provide the formulae for these algorithms in Appendix A. Finally, recall that  $\mathbf{L} \in \mathbb{R}^{m \times n}$  with  $L_{ij} = \mathbb{E}[Y_{ij}]$  and that  $\Omega$  denotes the set of observed action-context pairs.

**Local Approaches to Matrix Completion.** We call a causal prediction approach *local* if it predicts individual entries  $L_{ij}$ , rather than predicting the entire matrix  $\mathbf{L}$  in one go. Many local approaches predict  $L_{ij}$  via some weighted average  $\sum_{(k,\ell) \in \Omega} w_{k\ell} Y_{k\ell}$ , with different choices for  $w_{k\ell}$  giving rise to different methods. For example, in *Collaborative Filtering* (CF) (Schafer et al., 2007; Linden et al., 2003) and nearest-neighbors approaches (Dwivedi et al., 2022), the weight  $w_{k\ell}$  depends on some measure of data-dependent similarity (e.g. cosine similarity) between actions  $i, k$  and contexts  $j, \ell$ .

Very simple examples of local, weighting-based estimators include the *mean-over-contexts* estimator (which has  $w_{i'j'} \propto \mathbb{1}_{i'=i}$  with the weights summing to one), and the similarly-defined *mean-over-actions* estimator. The estimators are appropriate for models with *only* fixed action effects and fixed context effects, respectively. To handle a two-way fixed effect model (one with both fixed action and context effects), the mean-over-actions and mean-over-contexts estimators can be combined into the *Fixed Effects* (FE) estimator. More precisely, this estimator corresponds to taking an average of the fixed action effect estimators presented in Squires et al. (2022).

Connections between causal prediction and matrix completion are explored in a number of works. In particular, the widely-used *Synthetic Controls* method predicts outcomes for treated units under the counterfactual setting where they received no treatment (Abadie and Gardeazabal, 2003; Doudchenko and Imbens, 2016; Bai and Ng, 2019). Agarwal et al. (2020) generalized this idea to predict counterfactuals under treatment with the *Synthetic Interventions* (SI) estimator, which has also been connected to causal matrix/tensor completion (Agarwal et al., 2021; Squires et al., 2022).

**Global Approaches to Matrix Completion.** Other approaches, such as *nuclear norm minimization* (NNM) (Candes and Plan, 2009), predict the entire matrix  $\mathbf{L}$  in one go, often by casting matrix

Algorithm	Model Class	Symmetric	Dependence on $ \Omega $
Mean-Over-Contexts	Fixed Action Effect Model	$\times$	Low
Mean-Over-Actions	Fixed Context Effect Model	$\times$	Low
Fixed Effects (FE)	Two-way Fixed Effect Model	$\checkmark$	Low
Collaborative Filtering (CF)	Mixture Model	$\times$	Medium
Synthetic Interventions (SI)	Latent Factor Model	$\checkmark^*$	Medium
SI-mean-contexts	Latent Factor Model	$\times$	Low
SI-FE	Latent Factor Model	$\checkmark^*$	Low
Nuclear Norm Min. (NNM)	Low Rank	$\checkmark$	High
NNM-FE	Low Rank	$\checkmark$	Low

Table 1: **Summary of the algorithms.** *Model Class* indicates that the algorithm is consistent when the true data-generating process belongs to that class, e.g., Synthetic Interventions is consistent under a factor model. *Dependence on  $|\Omega|$*  (the number of observed entries) is a qualitative judgement based on the performance results shown in Section 5.

\*SI is symmetric for matrices but it is non-symmetric for third-order tensors.

completion as an optimization problem and developing fast optimizers (Cai et al., 2008; Candes and Recht, 2008; Mazumder et al., 2010). In fact, there are also global approaches that take advantage of latent factor models (Hastie et al., 2014; Jain et al., 2012). There is a vast literature studying the optimality properties of NNM and its variations (Candes and Tao, 2009; Recht et al., 2010; Zhang and Aeron, 2015). However, these results rely on the assumption of uniformly at random observations, which is usually not satisfied in real datasets, especially for biological applications (see Section 5). Thus, we might prefer to use *local approaches* when the missingness pattern is far from uniform.

Among global approaches, the one most related to the present work is an extension of NNM by Athey et al. (2021). This approach excludes fixed effects from regularization, hence we call it *NNM-FE*. Roughly speaking, if we let  $\hat{\mathbf{Y}}^{\text{FE}}$  denote the fixed effect estimator for  $\mathbf{Y}$ , NNM-FE is similar to using NNM on the matrix  $\mathbf{Y} - \hat{\mathbf{Y}}^{\text{FE}}$ . In Section 4, we use the same idea to improve the SI estimator when the model class includes fixed effects.

**Causal Prediction in Biological Applications.** We use the DepMap PRISM dataset (Corsello et al., 2020), which measures the viability score of drug and cell line combinations. Radhakrishnan et al. (2022) used additional information from the Connectivity Map (CMAP) dataset (Corsello et al., 2020) to predict the viability scores from DepMap, going beyond the index-only setting considered here. In the index-only setting, Squires et al. (2022) used Synthetic Interventions for the CMAP dataset while (Hodos et al., 2018) used a nearest-neighbor approach.

### 3. Background and Setup

We begin by reviewing relevant background on structural causal models (Peters et al., 2017), and then use the concepts to define the model class that we consider for causal prediction.

### 3.1. Background

**Definition 1 (Structural Causal Model (SCM))** A structural causal model (SCM) is defined by a tuple  $\mathcal{S} = (\mathbf{S}, \mathbb{P}^{\mathcal{E}})$ , where  $\mathbf{S} = (S_1, \dots, S_q)$  is an indexed set of  $q$  causal mechanisms

$$S_k : Z_k = f_k(\text{pa}(Z_k), \varepsilon_k), \quad k = 1, \dots, q.$$

Here,  $\text{pa}(Z_k) \subseteq \{Z_1, \dots, Z_q\} \setminus \{Z_k\}$  are called the parents of  $Z_k$ , and  $\mathbb{P}^{\mathcal{E}} = \mathbb{P}^{\varepsilon_1, \dots, \varepsilon_q}$  is the joint distribution of the exogenous noise variables  $\varepsilon_1, \dots, \varepsilon_q$ . The causal graph  $\mathcal{G}$  of an SCM has a directed edge  $Z_l \rightarrow Z_k$  for all  $k$  and for all  $Z_l \in \text{pa}(Z_k)$ . We assume that  $\mathcal{G}$  is acyclic, i.e., it is a DAG.

Unless otherwise noted, we assume that  $\mathbb{P}^{\mathcal{E}}$  is a product distribution, i.e., that the noise variables  $\varepsilon_1, \dots, \varepsilon_q$  are jointly independent. An SCM is called *Gaussian* if  $\varepsilon_k \sim \mathcal{N}(0, \sigma_k^2)$  with  $\sigma_k > 0 \forall k$ . It is called *linear* if all causal mechanisms  $f_k$  are linear. In particular, in a linear SCM, the causal mechanisms are defined by some parameters  $B_{l,k} \in \mathbb{R}$  and they can be written as

$$Z_k = \sum_{Z_l \in \text{pa}(Z_k)} B_{l,k} Z_l + \varepsilon_k \quad \text{for all } k = 1, \dots, q \quad (1)$$

We now define interventions and counterfactuals, mostly following the notation of [Peters et al. \(2017\)](#). Fix an SCM  $\mathcal{S} = (\mathbf{S}, \mathbb{P}^{\mathcal{E}})$  over nodes  $\mathbf{Z}$ . A (soft) *intervention*<sup>2</sup>  $I$  is defined by a set  $T(I) \subseteq \mathbf{Z}$  of *intervention targets*, and an indexed set  $\{f_k^I\}_{k \in T(I)}$  of *interventional causal mechanisms*, where  $f_k$  is generally allowed to be a function of  $\text{pa}(Z_k)$  and  $\varepsilon_k$ . Given an intervention  $I$ , we define the *interventional SCM* as  $\mathcal{S}_I = (\mathbf{S}^I, \mathbb{P}^{\mathcal{E}})$ , where  $(\mathbf{S}^I)_k = \mathbf{S}_k$  if  $k \in T(I)$  and  $(\mathbf{S}^I)_k = \mathbf{S}_k$  otherwise.

On the other hand, letting  $\mathbf{C} \subseteq \mathbf{Z}$  and conditioning on  $\mathbf{C} = \mathbf{c}$ , we define the *counterfactual SCM* as  $\mathcal{S}_{\mathbf{C}=\mathbf{c}} = (\mathbf{S}, \mathbb{P}^{\mathcal{E}|\mathbf{C}=\mathbf{c}})$ . To combine counterfactuals and interventions (in that order), we define  $\mathcal{S}_{[\mathbf{C}=\mathbf{c}, I]} = (\mathbf{S}^I, \mathbb{P}^{\mathcal{E}|\mathbf{C}=\mathbf{c}})$ . This SCM entails a new joint distribution over  $\mathbf{Z}$ , which we denote by  $\mathbb{P}(\mathbf{Z} \mid [\mathbf{C} = \mathbf{c}, I])$ , or in the special case of a do-intervention setting  $\mathbf{A} \subset \mathbf{Z}$  to  $\mathbf{a}$ ,  $\mathbb{P}(\mathbf{Z} \mid [\mathbf{C} = \mathbf{c}, \text{do}(\mathbf{A} = \mathbf{a})])$ .

### 3.2. An SCM-based Model Class for Causal Prediction

We give general results for the setting where we observe a vector-valued outcome  $\mathbf{Y}_{ij} \in \mathbb{R}^p$  for each action-context pair  $(i, j)$ . This gives rise to a partially observed 3-order tensor  $\mathbf{Y} \in \mathbb{R}^{m \times n \times p}$  where the rows (first index) correspond to actions and the columns (second index) correspond to contexts. Let  $\Omega$  denote the set of indices  $(i, j)$  for which we have observed  $\mathbf{Y}_{ij} \in \mathbb{R}^p$ . Our ultimate goal is to impute the missing entries of  $\mathbb{E}[\mathbf{Y}]$ . In the special case  $p = 1$ ,  $\mathbf{Y}$  reduces to an  $m \times n$  matrix, and tensor completion reduces to matrix completion.

We assume that there is an underlying SCM over some  $\mathbf{Z} = (Z_1, \dots, Z_q)$  that defines our system (each possible context). Let  $I_A$  and  $I_C$  be index variables that define the action and context, respectively. We add them to our SCM as follows. First, we observe  $I_C$  before applying any action. As  $I_C$  defines a context, it is defined as a function of  $\mathbf{Z}$ , so it is downstream from every node. Then, conditioning on  $I_C = j$  corresponds to conditioning on some  $\mathbf{Z}_{\mathcal{C}_j}$  where  $\mathcal{C}_j \subset [q] := \{1, \dots, q\}$ .

2. The terminology for different classes of interventions is fairly inconsistent. Soft interventions have many other names, e.g., *parametric interventions* ([Eberhardt and Scheines, 2007](#)) or *mechanism shifts* ([Tian and Pearl, 2001](#)). Our definition for soft interventions contains *perfect (hard) interventions* and *do-interventions* as special cases, see ([Squires and Uhler, 2022](#)) for additional terms for these classes.






Figure 3: **DAG defining our model class.** Shaded nodes are observed, while unshaded nodes are unobserved. Data is generated by conditioning on the context index  $I_C = j$  (indicated by the double-edge for the node  $I_C$ ) followed by intervening to set the action index  $I_A = i$  (indicated by the square node  $I_A$ ). The context may potentially depend on any subset of  $\mathbf{Z}$  and the action may potentially affect any subset of  $\mathbf{Z}$ . The exogenous noise terms  $\epsilon_1, \epsilon_2, \dots, \epsilon_q$  are shared between the pre-interventional and post-interventional SCMs. Here, the observed outcome is  $\mathbf{Y}_{ij} \sim \mathbb{P}(\mathbf{Z}_{\mathcal{X}}(i) \mid I_C = j)$  for  $\mathcal{X} = \{2, 4\}$ .

On the other hand, the index  $I_A$  can be thought of as an *instrumental variable* (Newey and Powell, 2003) or a *regime indicator* (Dawid, 2021), which encodes the (unknown) intervention that each action applies. In particular, using action  $i$  corresponds to setting  $I_A = i$ , which induces an intervention on some set of variables  $\mathbf{Z}_{\mathcal{A}_i}$ ,  $\mathcal{A}_i \subset [q]$ . Together, the intervention on  $I_A$  and conditioning on  $I_C$  gives rise a new SCM, and defines a new set of variables  $\mathbf{Z}(i)$  that represent the counterfactual state of the system.<sup>3</sup> Finally, we only observe a subset of  $\mathbf{Z}(i)$ , which we denote with some indices  $\mathcal{X} \subset [q]$ ,  $|\mathcal{X}| = p$ , and the outcome  $\mathbf{Y}_{ij}$  is drawn from the distribution  $\mathbb{P}(\mathbf{Z}_{\mathcal{X}}(i) \mid I_C = j)$ . This definition of our model class is summarized in Fig. 3. Altogether, the expected value of the  $(i, j)$ -th entry of  $\mathbf{L}$  is

$$\mathbb{E}(\mathbf{Y}_{ij}) = \mathbb{E}(\mathbf{Z}_{\mathcal{X}}(i) \mid I_C = j) \quad (2)$$

We clarify our setup with a simple example which uses only  $\text{do}$ -interventions.

**Example 1** Let  $Z_1 \rightarrow Z_2 \rightarrow Z_3$  follow an SCM  $\mathcal{S} = (\mathbf{S}, \mathbb{P}^{\mathcal{E}})$ . Assume that context is defined solely based on the value of the marker gene  $Z_3$ , e.g.,  $I_C = 1$  if and only if  $Z_3 = c_1$  and  $I_C = 2$  if and only if  $Z_3 = c_2$ . This gives us a new distribution on the noise,  $\mathbb{P}^{\mathcal{E} \mid Z_3 = c_j}$  for each possible  $c_j$ . Assume also that each action corresponds to a  $\text{do}$ -intervention on the gene  $Z_1$ , e.g. if  $I_A = 1$ , then we

3. This notation is inspired by Single World Intervention Graphs (SWIGs) (Richardson and Robins, 2013). Indeed, we would obtain a SWIG after performing  $\text{do}$ -interventions. However, we also consider soft interventions. We could write  $I_A = \emptyset$  for defining the control state, i.e. no action. In that case  $\mathbf{Z}(\emptyset) = \mathbf{Z}$ .

have performed the intervention  $d\circ(Z_1 = a_1)$ , if  $I_A = 2$ , then we have performed the intervention  $d\circ(Z_1 = a_2)$ . So for  $I_A = i$  we have a new set of structural equations  $S^{d\circ(Z_1=a_i)}$ . Therefore, the SCM after conditioning and intervening is  $S_{[Z_3=c_j, d\circ(Z_1=a_i)]} = (\mathbf{S}^{d\circ(Z_1=a_i)}, \mathbb{P}^{\mathcal{E}|Z_3=c_j})$ . Finally, assume that we use an assay which measures only the gene  $Z_2$ . Then our matrix satisfies

$$\mathbb{E}(\mathbf{Y}_{ij}) = \mathbb{E}(Z_2 \mid [Z_3 = c_j, d\circ(Z_1 = a_i)])$$

where we use the notation  $[Z_3 = c_j, d\circ(Z_1 = a_i)]$  to emphasize that the order of the terms is important: the intervention is considered in the model obtained after conditioning. Using our notation, shown in Equation (2), we would write  $\mathbb{E}(\mathbf{Y}_{ij}) = \mathbb{E}(Z_2(a_i) \mid Z_3 = c_j)$ .

### 3.3. Synthetic Interventions

The Synthetic Interventions (SI) estimator (Agarwal et al., 2020) is a local causal prediction approach for latent factor models. To predict  $\mathbf{Y}_{ij}$  for  $(i, j) \notin \Omega$ , SI follows the following procedure: (1) take the set of columns  $\mathcal{C}(i)$  for which we have observed row  $i$ , (2) take all rows for which we have observed columns  $\mathcal{C}(i) \cup \{j\}$ , (3) fit a linear regression on the available data defined by these subsets, and (4) use the linear regression coefficient to predict  $\mathbf{Y}_{ij}$  as a linear combination of  $\{\mathbf{Y}_{i\ell}\}_{\ell \in \mathcal{C}(i)}$ . Note that, analogously, we could use SI regressing along columns. See Appendix A for more details.

One of the main advantages of using SI is that, in contrast of NNM, it does not require significant assumptions on the missingness structure of the data. Instead, to prove theoretical guarantees, e.g. finite-sample consistency (Agarwal et al., 2020), the data should be generated from a latent factor model, and also satisfy a *linear span inclusion* assumption, namely that  $\mathbf{v}_j \in \text{span}(\mathbf{v}_k : k \in \mathcal{C}(i))$ . Intuitively, if  $\mathcal{C}(i)$  is sufficiently large in relation to the rank of our factor model, SI provides a consistent estimator.

## 4. Theoretical Results

In this section, we establish expressions for the expected value of the entries of our tensor/matrix  $\mathbf{Y}$ . To connect these expressions to specific tensor completion approaches, we will call an estimator *consistent* for an expression if it returns  $\mathbb{E}[\mathbf{Y}_{ij}]$  when given as input the expected values for each observed entry, i.e. the values  $\mathbb{E}[\mathbf{Y}_{uv}]$  for  $(u, v) \in \Omega$ . This corresponds to taking a limit, when each observed entry is the average of  $K$  independent samples of  $\mathbf{Y}_{uv} \sim \mathbb{P}_{I_C=v; d\circ(I_A=u)}$ , and we let  $K \rightarrow \infty$ . To simplify notation, for any  $\mathbf{v} \in \mathbb{R}^d$ , we define  $\text{aug}(\mathbf{v}) = (1, v_1, \dots, v_d)^\top \in \mathbb{R}^{d+1}$  as the vector given by prepending a 1 to  $\mathbf{v}$ . The proofs of the following results can be found in Appendix B.

**Assumption 1** *The entries of  $\mathbf{Y}$  come from a linear SCM over a set of variables  $\mathbf{Z} = (Z_1, \dots, Z_q)$  as in Equation (1). That is, the  $(i, j)$ -th entry corresponds to the counterfactual when conditioning on  $I_C = j$  and then intervening to set  $I_A = i$ , as in Equation (2).*

**Theorem 1** *Under Assumption 1, we have*

$$\mathbb{E}(\mathbf{Y}_{ij}) = \mathbb{E}(\mathbf{Z}_{\mathcal{X}}(i) \mid I_C = j) = U_i \mathbf{v}_i + U'_i \mathbf{w}_j.$$

for some  $U_i, U'_i \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{Z}|}$ ,  $\mathbf{v}_i \in \mathbb{R}^{|\mathcal{Z}|}$  depending on the action index  $i$ , and some  $\mathbf{w}_j \in \mathbb{R}^{|\mathcal{Z}|}$  depending on the context index  $j$ .



**Proof [Sketch]** For a linear SCM, we have  $\mathbf{Z} = (I - B)^{-1}\mathcal{E}$  with  $B \in \mathbb{R}^{q \times q}$  and  $\mathcal{E} = (\varepsilon_1, \dots, \varepsilon_q)$ . Conditioning on  $I_C = j$  gives a posterior  $\mathbb{P}^{\mathcal{E}|I_C=j}$  over the exogenous noise. Intervening gives a new matrix  $B_i$  and new exogenous noise  $\tilde{\mathcal{E}}_{\mathcal{A}_i}$ , where  $\mathcal{A}_i$  are the intervention targets when  $I_A = j$ . Defining  $\mathcal{E}' = \tilde{\mathcal{E}}_{\mathcal{A}_i} \cup \mathcal{E}_{\overline{\mathcal{A}_i}}$ , then  $\mathbf{Z}(i)$  equals  $(I - B_i)^{-1}\mathcal{E}'$  in distribution. Thus,  $\mathbb{E}(\mathbf{Z}(i) | I_C = j) = U_i \mathbf{v}_i + U'_i \mathbf{w}_j$  for  $U_i = (I - B_i)^{-1}$  and  $U'_i$  a masked version of  $U_i$ , with  $\mathbf{v}_i$  and  $\mathbf{w}_j$  coming from  $\mathbb{E}[\mathcal{E}']$ . ■

Theorem 1 shows that  $\mathbb{E}[\mathbf{Y}]$  follows a factor model with a fixed action effect. The fixed effect can be folded into the factor model as follows:

$$U_i \mathbf{v}_i + U'_i \mathbf{w}_j = \underbrace{\left( \begin{array}{c|c} U_i \mathbf{v}_i & U'_i \end{array} \right)}_{\mathbb{R}^{|\mathcal{X}| \times (1+|\mathcal{Z}|)}} \text{aug}(\mathbf{w}_j) \quad (3)$$

Thus, Synthetic Interventions will be consistent for completing the tensor. However, we propose an alternative to deal with the fixed effect.

#### 4.1. Theoretically-motivated extensions of Synthetic Interventions

The factor model in Equation (3) has more structure than a general rank  $|\mathcal{Z}| + 1$  factor model. We have a fixed effect along rows plus a factor model of rank  $|\mathcal{Z}|$ . The best approximation of a one-way fixed effect is given by the mean-over-contexts estimator (see Appendix C). Therefore, following the same idea from the NNM-FE [Athey et al. \(2021\)](#), we propose the following modification to SI: (1) subtract the fixed effect, i.e., let  $\mathbf{D} = \mathbf{Y} - \hat{\mathbf{Y}}^{\text{mean-over-contexts}}$ , (2) run SI on this matrix and obtain  $\hat{\mathbf{D}}^{\text{SI}}$ , (3) return the estimate  $\hat{\mathbf{Y}}^{\text{mean-over-contexts}} + \hat{\mathbf{D}}^{\text{SI}}$ . Intuitively, by removing the exact fixed effect, our factor model would have rank  $|\mathcal{Z}|$  instead of  $|\mathcal{Z}| + 1$ , so it would be easier to satisfy the linear span inclusion assumption.

**Tensor case.** The non-symmetry in the factor model from Equation (3) is even more relevant in the tensor case, i.e.  $|\mathcal{X}| > 1$  (we have a matrix that depends on  $i$  and a vector that depends on  $j$ ). Therefore, it is more reasonable to regress along contexts than along actions. To show how SI is non-symmetric for tensors, in Appendix D, we provide a detailed example for a two-node case where we observe both  $Z_1$  and  $Z_2$ . Following our setup, we provide a  $4 \times 4 \times 2$  tensor that can only be completed if we use SI within columns (contexts). In particular, for using SI within rows we would need more data to satisfy the linear span inclusion assumption.

#### 4.2. Exploiting Additional Structure in the Model Class

One drawback about the factor model from Equation (3) is that it has rank  $|\mathcal{Z}| + 1$ , which might be considerably large. We now consider two additional assumptions which impose additional structure on how  $I_A$  and  $I_C$  interact with the SCM over  $\mathbf{Z}$ . First, we consider models where the all contexts are defined by some common set of system variables.

**Assumption 2** *There is some  $\mathcal{C} \subset [q]$  such that  $\mathcal{C}_j = \mathcal{C}$  for all  $j \in [n]$ .*

In our setting, where we observe only a very coarse-grained context index, such as cell type, it is reasonable to suppose that (at least approximately) the index can be determined from only a small set of *context markers*  $\mathcal{C}$ . For example, there may exist some small set of marker genes that define cell states. Assumption 2 is important to consider, since it can have a substantial effect on identifiability and estimation by reducing the rank of our factor model to  $|\mathcal{C}|$ .

**Proposition 2** *Let Assumptions 1 and 2 hold, and assume that the underlying SCM is Gaussian. Then the latent factor model from Theorem 1 is*

$$U_i' \mathbf{w}_j = U_i' W \mathbf{c}_j$$

where  $W$  is a constant matrix and  $\mathbf{c}_j$  is the value that context markers take to define context  $j$ . Therefore, by defining  $U_i'' = U_i' W \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{C}|}$  we obtain a latent factor model of rank  $|\mathcal{C}|$ .

**Assumption 3** *There is some  $\mathcal{A} \subset [q]$  such that  $A_i = \mathcal{A}$  for all  $i \in [m]$ .*

Assumption 3 indicates that all actions change the SCM in a similar way. Although this may not be realistic in most settings, we might have this setup for a particular sub-matrix of our whole matrix  $\mathbf{Y}$ . For example, Assumption 3 may hold for drugs which use the same mechanism of actions, or if the different actions correspond to the same drug over different dosage concentrations. In this case, the latent factor model reduces to a two-way fixed effects model.

**Corollary 3** *Under Assumptions 1 and 3 we have*

$$\mathbb{E}(\mathbf{Z}_{\mathcal{X}}(i) \mid I_C = j) = U \mathbf{v}_i + U' \mathbf{w}_j.$$

so Fixed Effect (FE) is a consistent estimator for completing  $\mathbf{Y}$ .

In our previous results, we have considered only the case of a linear SCM. This can be extended in the simplified case where (1) all interventions occur on the same set of nodes and (2) all contexts use the same context-defining nodes. While linearity may appear to be a strong assumption, it is closely related to the low-rank assumption, which is essential in matrix completion. In Appendix E, we study some examples of non-linear SCMs. In particular, when considering polynomials, we observe how the rank of the LFM increases as the degree of the polynomials increases. These results provide evidence that some form of linearity or low-degree polynomial assumption may be required to justify a connection between factor models and causal models.

Finally, in real applications, we need to test the validity of our assumptions, i.e. we need to test whether our tensor comes from a certain graphical model. So far, we have only exploited the structure that the SCM induces on *expectations* of  $\mathbf{Y}$ . However, we can also use the structure to obtain fine-grained implications regarding the variance of our observations. This structure on the variance is helpful for inference and for hypothesis testing. We demonstrate this idea in Appendix F, where we propose a hypothesis test for the fixed effects model implied by Corollary 3. In particular, we test Corollary 3 implies homoscedastic errors within a row. Thus, we develop a hypothesis test for homoscedasticity, and then a particular test for fixed effects which reduces to a Welch's test. Similarly, for the latent factor model (LFM) case from Prop. 2, we can test homoscedasticity within columns, then use the test for LFMs proposed by Agarwal et al. (2020). In Appendix F, we demonstrate the performance of these tests on synthetic data.

## 5. Empirical results

We work with the PRISM Repurposing dataset (Corsello et al., 2020). The entries of  $\mathbf{Y}$  measure a viability score, which indicates how lethal a drug is for a given cell line (i.e., negative viability means that a large proportion of cells die). Since we consider the index-only setting, we use only the

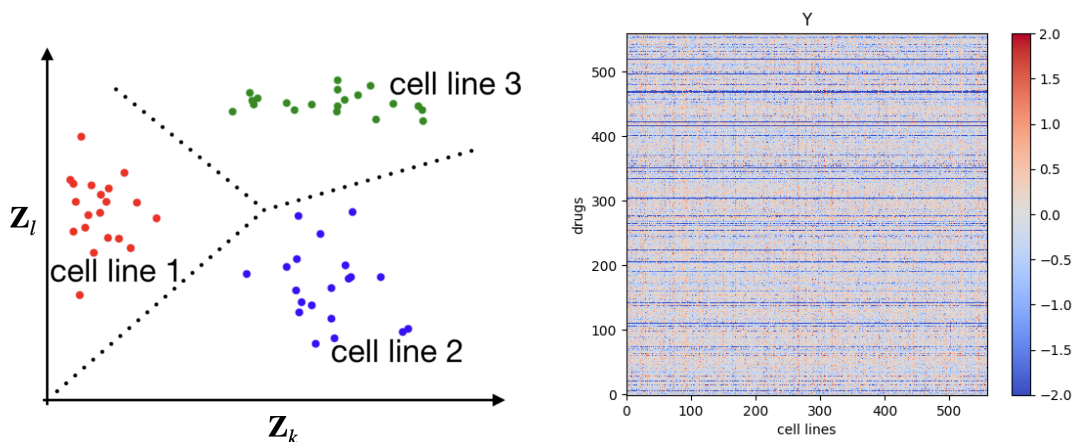


Figure 4: **(Left) Translating the PRISM data to our model class.** Letting  $\mathcal{C} = \{k, l\}$  indicates that cell state is defined solely in terms of  $Z_k$  and  $Z_l$ . **(Right) Matrix of viability scores.** Each entry represents the viability score for a drug-cell line pair. Negative viability indicates cell death. Viability scores are normalized for visualization purposes.

observed entries of this matrix to make our predictions. To relate the PRISM data into the model class introduced in Section 4, we think of the initial cell states in terms of Fig. 4 (left). In particular, we assume that there is some latent space (defined by the context variables) where we can perfectly distinguish different cell lines, e.g. the expression levels of some marker genes.

Arranging drugs in rows and cell lines in columns, as in Fig. 1 from Section 1, we obtain a  $4,686 \times 568$  matrix. To check the symmetry of an algorithm, we want the same number of rows as columns, so that completing along one direction (e.g. columns) is not better simply because there is more data to make predictions. Thus, we run the experiments on a square submatrix to avoid of this bias. For the same reason, we use symmetric missing data patterns, as described in Section 5.1. We compare the performance of the algorithms in terms of the  $R^2$  score. The baseline used in the denominator of the  $R^2$  is the average over the missing data. As an example, for the matrix shown in Fig. 4 (right) and a missing pattern from Fig. 5, the baseline MSE is approximately 0.87. Getting high  $R^2$  is a difficult task because the matrix is relatively noisy (see Appendix G.1).




Figure 5: **Missing data patterns used in our experiments.** Observed entries are denoted with black. For each missing entry we have the same number of observations along rows and columns.

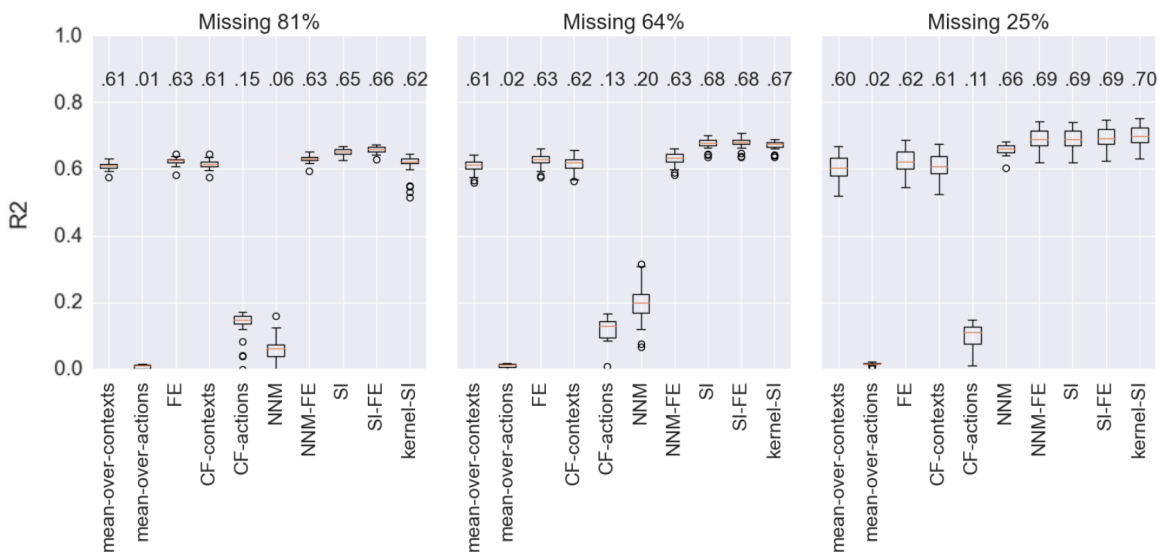


Figure 6: **Variants of Synthetic Interventions have the best performance on matrix completion for the PRISM dataset.** Using the missing data patterns are ones shown in Fig. 5, we test the performance of each algorithm on 20 different shuffles of rows and columns. The results are shown using boxplots. The numbers at the top denote the median (red lines).

### 5.1. About the missing data pattern

In the PRISM dataset, we have observed all the matrix. As a consequence, we can decide the missing data pattern to test the performance of the matrix completion approaches. But which missing data pattern is more appropriate for biological data?

As we have seen in Section 2, Nuclear Norm Minimization algorithms are near-optimal when we have uniformly at random missing entries. However, in biology applications this is usually not the case. Instead, it is more common to have a small portion of cell lines that has been tested against many of the drugs, but for the vast majority of cell lines we have only run a few experiments. For example, in Appendix G.2 we can see the missing data pattern for the CMAP dataset.

For this reason, we test the algorithms on missing patterns as the ones from Fig. 5. Here, for each missing entry  $(i, j)$ , the number of non-missing entries (black squares) in the  $i$ -th row is equal to the number of non-missing entries in the  $j$ -th column. We run the experiments using different data patterns to see the effect of having less/more available data.

In Appendix G.3 we consider an alternative missing data pattern where we do not have a constant number of observations for each missing entry. We can observe similar results also in this setting.

### 5.2. Performance of matrix completion algorithms

In this section we show the performance of the matrix completion algorithms presented before. In particular, we are interested in analyzing how this performance depends on the amount of observed data. An extensive analysis of this dependence can be found in Appendix G.4.

In Fig. 6 we see that mean-over-contexts is a strong baseline and that Synthetic Interventions (and its variants) outperform all the other approaches. NNM performs competitively when we observe a sufficiently large amount of entries (right-hand side plot) but it performs quite poorly otherwise.

These results reinforce our theoretical findings. In Section 4 we see how, assuming an underlying linear SCM, the counterfactual quantity of our interest leads to a particular factor model, implying consistency of the Synthetic Interventions estimator. For further analysis, in Appendix G.5 we report results for other variations of these approaches. In Appendix G.6 we run similar experiments on a particular submatrix.

## 6. Discussion

In this paper, we generalize the SCM-based model class introduced by Squires et al. (2022) to allow for situations where the measured outcome is defined as a counterfactual. We showed that latent factor models (LFMs) naturally arise out of this model class, though with an additional term for fixed action effects. This demonstrates the fundamentally distinctive nature of *causal* matrix completion, where there is a difference between conditioning (completing within columns) and intervening (completing within rows). As a practical consequence, this led us to propose a simple extension of the Synthetic Interventions (SI) estimator which includes a fixed effect term. These model-inspired causal matrix completion approaches work considerably well, especially in the low-data regime.

**Limitations and Future Work.** In this paper, we largely focused on linear models, and only sought to develop *consistent* estimators for causal predictions. An important next step is to analyze the noisy case, focusing on statistical aspects such as sample complexity and efficient inference.

A key limitation of our work is that we only considered *index-only* causal prediction problems. In many applications, additional data is available for actions and/or contexts, e.g., the molecular structure of a drug would be highly relevant to predicting its effect. We expect that the current work will be a useful conceptual foundation for developing model-based approaches to these causal prediction problems, especially when developing further connections between SCMs and latent factor models. It would be especially interesting to explore connections with recent works for intervention extrapolation, such as Agarwal et al. (2023) and Saengkyongam et al. (2023).

Finally, while our work was mainly motivated by a biological problem, the causal prediction problem is very general, and it would be interesting to apply our framework in other settings.

## Acknowledgments

This work was supported in part by funding from the Eric and Wendy Schmidt Center at the Broad Institute of MIT and Harvard.

Á. Ribot acknowledges support by the mobility grants program of Centre de Formació Interdisciplinària Superior (CFIS) - Universitat Politècnica de Catalunya (UPC), the MIT Simons MMLS Foundation research grant (6941629), and a MOBINT-MIF grant.

C. Squires and C. Uhler acknowledge support by the NSF TRIPODS program (DMS-2022448), NCCIH/NIH (1DP2AT012345), ONR (N00014-22-1-2116), the United States Department of Energy (DOE), Office of Advanced Scientific Computing Research (ASCR), via the M2dt MMICC center (DE-SC0023187), AstraZeneca, the Eric and Wendy Schmidt Center at the Broad Institute, and a Simons Investigator Award.

## References

- Alberto Abadie and Javier Gardeazabal. The economic costs of conflict: A case study of the basque country. *American Economic Review*, 93(1):113–132, March 2003. doi: 10.1257/000282803321455188. URL <https://www.aeaweb.org/articles?id=10.1257/000282803321455188>.
- Abhineet Agarwal, Anish Agarwal, and Suhas Vijaykumar. Synthetic combinations: A causal inference framework for combinatorial interventions, 2023.
- Anish Agarwal, Devavrat Shah, and Dennis Shen. Synthetic a/b testing using synthetic interventions, 2020. URL <https://arxiv.org/abs/2006.07691>.
- Anish Agarwal, Munther Dahleh, Devavrat Shah, and Dennis Shen. Causal matrix completion, 2021.
- Susan Athey, Mohsen Bayati, Nikolay Doudchenko, Guido Imbens, and Khashayar Khosravi. Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 116(536):1716–1730, may 2021. doi: 10.1080/01621459.2021.1891924. URL <https://doi.org/10.1080%2F01621459.2021.1891924>.
- Jushan Bai and Serena Ng. Matrix completion, counterfactuals, and factor analysis of missing data, 2019. URL <https://arxiv.org/abs/1910.06677>.
- Jian-Feng Cai, Emmanuel J. Candes, and Zuowei Shen. A singular value thresholding algorithm for matrix completion, 2008.
- Emmanuel J. Candes and Yaniv Plan. Matrix completion with noise, 2009.
- Emmanuel J. Candes and Benjamin Recht. Exact matrix completion via convex optimization, 2008. URL <https://arxiv.org/abs/0805.4471>.
- Emmanuel J. Candes and Terence Tao. The power of convex relaxation: Near-optimal matrix completion, 2009. URL <https://arxiv.org/abs/0903.1476>.
- Daniel C Castro, Ian Walker, and Ben Glocker. Causality matters in medical imaging. *Nature Communications*, 11(1):3673, 2020.
- Krzysztof Chalupka, Pietro Perona, and Frederick Eberhardt. Visual causal feature learning. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 181–190, 2015.
- Steven M Corsello, Rohith T Nagari, Ryan D Spangler, Jordan Rossen, Mustafa Kocak, Jordan G Bryan, Ranad Humeidi, David Peck, Xiaoyun Wu, Andrew A Tang, et al. Discovering the anticancer potential of non-oncology drugs by systematic viability profiling. *Nature cancer*, 1(2): 235–248, 2020.
- Philip Dawid. Decision-theoretic foundations for statistical causality. *Journal of Causal Inference*, 9(1):39–77, 2021.



- Nikolay Doudchenko and Guido W Imbens. Balancing, regression, difference-in-differences and synthetic control methods: A synthesis. Technical report, National Bureau of Economic Research, 2016.
- Raaz Dwivedi, Katherine Tian, Sabina Tomkins, Predrag Klasnja, Susan Murphy, and Devavrat Shah. Doubly robust nearest neighbors in factor models, 2022.
- Frederick Eberhardt and Richard Scheines. Interventions and causal inference. *Philosophy of science*, 74(5):981–995, 2007.
- Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, September 1936.
- Amir Feder, Katherine A Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E Roberts, et al. Causal inference in natural language processing: Estimation, prediction, interpretation and beyond. *Transactions of the Association for Computational Linguistics*, 10:1138–1158, 2022.
- Trevor Hastie, Rahul Mazumder, Jason Lee, and Reza Zadeh. Matrix completion and low-rank svd via fast alternating least squares, 2014.
- Rachel Hodos, Ping Zhang, Hao-Chih Lee, Qiaonan Duan, Zichen Wang, Neil R Clark, Avi Ma’ayan, Fei Wang, Brian Kidd, Jianying Hu, et al. Cell-specific prediction and application of drug-induced gene expression profiles. In *PACIFIC SYMPOSIUM ON BIOCOMPUTING 2018: Proceedings of the Pacific Symposium*, pages 32–43. World Scientific, 2018.
- Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization, 2012.
- Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009. doi: 10.1109/MC.2009.263.
- G. Linden, B. Smith, and J. York. Amazon.com recommendations: item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80, 2003. doi: 10.1109/MIC.2003.1167344.
- Daniel Malinsky, Ilya Shpitser, and Thomas Richardson. A potential outcomes calculus for identifying conditional path-specific effects, 2019.
- Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, 11(80):2287–2322, 2010. URL <http://jmlr.org/papers/v11/mazumder10a.html>.
- Whitney K Newey and James L Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, 2003.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- Adityanarayanan Radhakrishnan, Max Ruiz Luyten, Neha Prasad, and Caroline Uhler. Transfer learning with kernel methods, 2022.

- Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, jan 2010. doi: 10.1137/070697835. URL <https://doi.org/10.1137%2F070697835>.
- Thomas S Richardson and James M Robins. Single world intervention graphs (swigs): A unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper*, 128(30):2013, 2013.
- Sorawit Saengkyongam, Elan Rosenfeld, Pradeep Ravikumar, Niklas Pfister, and Jonas Peters. Identifying representations for intervention extrapolation, 2023.
- J Ben Schafer, Dan Frankowski, Jon Herlocker, and Shilad Sen. Collaborative filtering recommender systems. *The adaptive web: methods and strategies of web personalization*, pages 291–324, 2007.
- Christopher Southan, Markus Sitzmann, and Sorel Muresan. Comparing the chemical structure and protein content of chembl, drugbank, human metabolome database and the therapeutic target database. *Molecular informatics*, 32:881–897, 12 2013. doi: 10.1002/minf.201300103.
- Chandler Squires and Caroline Uhler. Causal structure learning: A combinatorial perspective. *Foundations of Computational Mathematics*, pages 1–35, 2022.
- Chandler Squires, Dennis Shen, Anish Agarwal, Devavrat Shah, and Caroline Uhler. Causal imputation via synthetic interventions. In *Conference on Causal Learning and Reasoning*, pages 688–711. PMLR, 2022.
- Jin Tian and Judea Pearl. Causal discovery from changes. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pages 512–521, 2001.
- Madeleine Udell and Alex Townsend. Why are big data matrices approximately low rank?, 2017.
- You Wu, Qiao Liu, Yue Qiu, and Lei Xie. Deep learning prediction of chemical-induced dose-dependent and context-specific multiplex phenotype responses and its application to personalized alzheimer’s disease drug repurposing. *PLOS Computational Biology*, 18(8):e1010367, 2022.
- Zemin Zhang and Shuchin Aeron. Exact tensor completion using t-svd, 2015.

**Contents of Appendix**

<b>A</b>	<b>Summary of matrix completion approaches</b>	<b>18</b>
<b>B</b>	<b>Proofs</b>	<b>19</b>
<b>C</b>	<b>Best approach for estimating a one-way fixed effect</b>	<b>21</b>
<b>D</b>	<b>SI is non-symmetric in the tensor case</b>	<b>21</b>
<b>E</b>	<b>Polynomial case</b>	<b>23</b>
E.1	Going beyond polynomials . . . . .	23
<b>F</b>	<b>Hypothesis tests</b>	<b>24</b>
F.1	Homoscedasticity within rows/columns . . . . .	24
F.2	Test for fixed effects . . . . .	25
F.3	Test for mean-over-contexts . . . . .	26
F.4	Test for Synthetic Interventions . . . . .	27
<b>G</b>	<b>Further analysis on PRISM Repurposing dataset</b>	<b>28</b>
G.1	SVD analysis . . . . .	28
G.2	Missing pattern in CMAP dataset . . . . .	29
G.3	An alternative missing data pattern . . . . .	30
G.4	Performance of the algorithms depending on the amount of data observed . . . . .	31
G.5	Performance of all the matrix completion algorithms considered . . . . .	33
G.6	Removal of "killer drugs" . . . . .	33

## Supplemental material

All code for running the experiments presented in this paper can be found at

<https://github.com/alvaro-ribot/causal-matrix-completion-PRISM>

## Appendix A. Summary of matrix completion approaches

In Section 2, we presented the different matrix completion approaches considered in this paper. For completeness, in this section, we provide formulae for such approaches.

We follow the notation used in Squires et al. (2022). Let  $\mathbf{Y} = (Y_{ij}) \in \mathbb{R}^{m \times n}$  be our matrix of interest. Let  $\Omega \subset [m] \times [n]$  be the set of observed entries and  $\mathcal{M} = [m] \times [n] \setminus \Omega$  be the set of missing entries. For a given row  $i$ , let  $\mathcal{C}(i) = \{j : (i, j) \in \Omega\}$  be the set of column indices  $j$  for which we have observed  $Y_{ij}$ . Similarly, for a column  $j$  define  $\mathcal{R}(j) = \{i : (i, j) \in \Omega\}$ . This notation is extended to sets of rows ( $\mathcal{R}$ ) and sets of columns ( $\mathcal{C}$ ):  $\mathcal{C}(\mathcal{R}) = \bigcap_{i \in \mathcal{R}} \mathcal{C}(i)$  and  $\mathcal{R}(\mathcal{C}) = \bigcap_{j \in \mathcal{C}} \mathcal{R}(j)$ . The following example clarifies our notation.

$$\boxed{\begin{array}{l} \begin{pmatrix} Y_{11} & Y_{12} & ? & ? & Y_{15} \\ ? & ? & Y_{23} & Y_{24} & ? \\ Y_{31} & ? & Y_{33} & ? & Y_{35} \end{pmatrix} \quad \begin{array}{l} \mathcal{C}(1) = \{1, 2, 5\} \\ \mathcal{C}(3) = \{1, 3, 5\} \end{array} \\ \implies \mathcal{C}(\{1, 3\}) = \{1, 5\} \end{array}}$$

- Mean-Over-Contexts:  $\hat{Y}_{ij}^{\text{mean-over-contexts}} = \frac{1}{|\mathcal{C}(i)|} \sum_{j' \in \mathcal{C}(i)} Y_{ij'}$ .
- Mean-Over-Actions:  $\hat{Y}_{ij}^{\text{mean-over-actions}} = \frac{1}{|\mathcal{R}(j)|} \sum_{i' \in \mathcal{R}(j)} Y_{i'j}$ .
- Fixed Effects (FE):  $\hat{Y}_{ij}^{\text{FE}} = \hat{Y}_{ij}^{\text{mean-over-contexts}} + \hat{Y}_{ij}^{\text{mean-over-actions}} - \frac{1}{|\mathcal{R}(j)||\mathcal{C}(i)|} \sum_{\substack{i' \in \mathcal{R}(j) \\ j' \in \mathcal{C}(i)}} Y_{i'j'}$ .
- Collaborative Filtering (CF):  $\hat{Y}_{ij}^{\text{CF-contexts}} = \frac{1}{\sum_{j' \in \mathcal{R}(i)} |\cos\text{-sim}(j', j)|} \sum_{j' \in \mathcal{R}(i)} \cos\text{-sim}(j', j) Y_{ij'}$ .
- Synthetic interventions (SI): Let  $\mathcal{R}_{\text{train}} = \mathcal{R}(\mathcal{C}(i) \cup \{j\})$ ,  $\mathbf{Y}_{\text{train}} = [Y_{i'j'} : i' \in \mathcal{R}_{\text{train}}, j' \in \mathcal{C}(i)]$ ,  $\mathbf{y}_{\text{output}} = [Y_{i'j} : i' \in \mathcal{R}_{\text{train}}]$ ,  $\mathbf{y}_{\text{test}} = [Y_{ij'} : j' \in \mathcal{C}(i)]$ . Then solve linear regression:

$$\hat{Y}_{ij}^{\text{SI}} = \mathbf{y}_{\text{test}}^\top \hat{\beta}, \quad \text{where } \hat{\beta} = \mathbf{Y}_{\text{train}}^\dagger \mathbf{y}_{\text{output}} \text{ and } \dagger \text{ denotes the pseudoinverse.}$$

- SI-FE: Let  $\mathbf{D} = \mathbf{Y} - \hat{\mathbf{Y}}^{\text{FE}}$ , then  $\hat{Y}_{ij}^{\text{SI-FE}} = \hat{Y}_{ij}^{\text{FE}} + \hat{D}_{ij}^{\text{SI}}$ .
- Nuclear Norm Minimization (NNM):  $\hat{\mathbf{Y}}^{\text{NNM}} = \arg \min_{\mathbf{L}} \frac{1}{2} \|\mathbf{P}_\Omega(\mathbf{Y} - \mathbf{L})\|_{\text{F}}^2 + \lambda \|\mathbf{L}\|_*$ , where  $(\mathbf{P}_\Omega(\mathbf{M}))_{ij} = M_{ij}$  if  $(i, j) \in \Omega$  and 0 otherwise.
- NNM-FE:  $\hat{\mathbf{Y}}^{\text{NNM-FE}} = \hat{\mathbf{L}} + \hat{\Gamma} \mathbf{1}_n^\top + \mathbf{1}_m \hat{\Delta}^\top$ , where  $\mathbf{1}_n = (1, \dots, 1) \in \mathbb{R}^n$  and

$$(\hat{\mathbf{L}}, \hat{\Gamma}, \hat{\Delta}) = \arg \min_{\mathbf{L}, \Gamma, \Delta} \frac{1}{|\Omega|} \|\mathbf{P}_\Omega(\mathbf{Y} - \mathbf{L} - \Gamma \mathbf{1}_n^\top - \mathbf{1}_m \Delta^\top)\|_{\text{F}}^2 + \lambda \|\mathbf{L}\|_*.$$

## Appendix B. Proofs

**Notation.** Recall that  $[q] := \{1, \dots, q\}$ . Given a tuple of indices  $\mathcal{K} = (k_1, \dots, k_p)$  where  $k_r \in [q]$  for all  $r = 1, \dots, p$ , we define the matrix  $\mathbb{I}_{\mathcal{K}}$  as follows:

$$\mathbb{I}_{\mathcal{K}} \in \mathbb{R}^{q \times q} \text{ s.t. } (\mathbb{I}_{\mathcal{K}})_{rs} = \begin{cases} 1 & \text{if } r = s \text{ and } r \in \mathcal{K} \\ 0 & \text{otherwise} \end{cases}$$

Furthermore, given a vector  $\mathbf{Z} \in \mathbb{R}^q$ , we define  $\mathbf{Z}_{\mathcal{K}} \in \mathbb{R}^{|\mathcal{K}|}$  such that  $(\mathbf{Z}_{\mathcal{K}})_l = \mathbf{Z}_{k_l}$ . Finally, given a matrix  $M \in \mathbb{R}^{q \times q}$ , we denote by  $[M]_{\mathcal{K}} \in \mathbb{R}^{|\mathcal{K}| \times q}$  the sub-matrix of  $M$  whose rows indices are in  $\mathcal{K}$ . To simplify notation, for any  $\mathbf{v} \in \mathbb{R}^d$ , we define  $\text{aug}(\mathbf{v}) = (1, v_1, \dots, v_d)^\top \in \mathbb{R}^{d+1}$  as the vector given by prepending a 1 to  $\mathbf{v}$ .

**Theorem 1** *Under Assumption 1, we have*

$$\mathbb{E}(\mathbf{Y}_{ij}) = \mathbb{E}(\mathbf{Z}_{\mathcal{X}}(i) \mid I_C = j) = U_i \mathbf{v}_i + U_i' \mathbf{w}_j.$$

for some  $U_i, U_i' \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{Z}|}$ ,  $\mathbf{v}_i \in \mathbb{R}^{|\mathcal{Z}|}$  depending on the action index  $i$ , and some  $\mathbf{w}_j \in \mathbb{R}^{|\mathcal{Z}|}$  depending on the context index  $j$ .

**Proof** Let  $\mathbf{Z} = (Z_1, \dots, Z_q)$  follow a linear SEM, i.e.

$$\mathbf{Z} = B\mathbf{Z} + \mathcal{E}$$

where  $B = (B_{ij}) \in \mathbb{R}^{q \times q}$  such that  $B_{ij} = 0$  whenever  $i \leq j$  and  $\mathcal{E} = (\varepsilon_1, \dots, \varepsilon_q)$ . Let  $I$  be the identity  $q \times q$  matrix. Since  $B$  is lower triangular,  $(I - B)$  is invertible and we have that

$$\mathbf{Z} = (I - B)^{-1} \mathcal{E}$$

When we condition on the value of the context index, i.e. on  $I_C = j$ , we get a new distribution for the noise,  $\mathbb{P}^{\mathcal{E} \mid I_C = j}$ . In particular,

$$\mathbf{w}_j := \mathbb{E}(\mathcal{E} \mid I_C = j) \in \mathbb{R}^q$$

is a function of  $j$ .

Now let  $\mathcal{A}_i \subseteq [q]$  be the target set of indices for the intervention with index  $i$ . For each  $k \in \mathcal{A}_i$ , the  $k$ -th row of  $B$  is changed (modifying the dependency from its parents). Let  $B_i$  be the weight matrix after the intervention<sup>4</sup>.

Moreover, for each  $k \in \mathcal{A}_i$ , we have a new noise variable  $\tilde{\varepsilon}_k$ , independent of  $\mathcal{E}$ . So, the expected value of these new variables depend on the intervention index  $i$  but not on the context index  $j$ . Let  $\tilde{\mathcal{E}}$  be a  $q$ -dimensional random vector such that its  $k$ -th entry is  $\tilde{\varepsilon}_k$  if  $k \in \mathcal{A}_i$  and 0 otherwise. Define

$$\mathbf{v}_i := \mathbb{E}(\tilde{\mathcal{E}}) \in \mathbb{R}^q, \quad \text{and note that } \mathbb{E}(\tilde{\mathcal{E}}) = \mathbb{E}(\tilde{\mathcal{E}} \mid I_C = j).$$

Therefore, the variables  $\mathbf{Z}(i)$  after the intervention satisfy the following SEM

$$\mathbf{Z}(i) = B_i \mathbf{Z}(i) + \tilde{\mathcal{E}} + \mathbb{I}_{\mathcal{A}_i} \mathcal{E} \implies \mathbf{Z}(i) = (I - B_i)^{-1} \left( \tilde{\mathcal{E}} + \mathbb{I}_{\mathcal{A}_i} \mathcal{E} \right)$$

4. In the particular case of do-interventions, we would remove these dependencies, so all the entries in the  $k$ -th row of  $B_i$  would be 0.

where  $\overline{\mathcal{A}_i} = [q] \setminus \mathcal{A}_i$ . Hence,

$$\mathbb{E}(\mathbf{Z}(i) \mid I_C = j) = (I - B_i)^{-1} \mathbb{E}(\tilde{\mathcal{E}}) + (I - B_i)^{-1} \mathbb{I}_{\overline{\mathcal{A}_i}} \mathbb{E}(\mathcal{E} \mid I_C = j)$$

Recall that we are interested only in  $\mathbf{Z}_{\mathcal{X}}$  (the variables we observe). Let

$$U_i = [(I - B_i)^{-1}]_{\mathcal{X}} \in \mathbb{R}^{|\mathcal{X}| \times q} \quad U'_i = U_i \mathbb{I}_{\overline{\mathcal{A}_i}} \in \mathbb{R}^{|\mathcal{X}| \times q}$$

Therefore, we have

$$\mathbb{E}(\mathbf{Z}_{\mathcal{X}}(i) \mid I_C = j) = U_i \mathbf{v}_i + U'_i \mathbf{w}_j$$

■

**Proposition 2** *Let Assumptions 1 and 2 hold, and assume that the underlying SCM is Gaussian. Then the latent factor model from Theorem 1 is*

$$U'_i \mathbf{w}_j = U'_i W \mathbf{c}_j$$

where  $W$  is a constant matrix and  $\mathbf{c}_j$  is the value that context markers take to define context  $j$ . Therefore, by defining  $U''_i = U'_i W \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{C}|}$  we obtain a latent factor model of rank  $|\mathcal{C}|$ .

**Proof** We have  $\mathcal{E} = (\varepsilon_1, \dots, \varepsilon_q) \sim \mathcal{N}(\mu_{\varepsilon} = \mathbf{0}, \Sigma_{\varepsilon} = \text{diag}(\sigma_i^2))$ , so

$$\Sigma := \text{Cov}(\mathbf{Z}) = (I - B)^{-1} \Sigma_{\varepsilon} (I - B^{\top})^{-1}$$

Let  $\mathbf{Z}_{\mathcal{C}}$  be the context markers variables and  $\overline{\mathcal{C}} = [q] \setminus \mathcal{C}$ . Using the Schur Complement we obtain

$$\mathbb{E}(\mathbf{Z}_{\overline{\mathcal{C}}} \mid \mathbf{Z}_{\mathcal{C}} = \mathbf{c}_j) = \Sigma_{\overline{\mathcal{C}}\mathcal{C}} \Sigma_{\mathcal{C}\mathcal{C}}^{-1} \mathbf{c}_j$$

Of course, we also have  $\mathbb{E}(\mathbf{Z}_{\mathcal{C}} \mid \mathbf{Z}_{\mathcal{C}} = \mathbf{c}_j) = \mathbf{c}_j$ , so  $\mathbb{E}(\mathbf{Z} \mid \mathbf{Z}_{\mathcal{C}} = \mathbf{c}_j) = M \mathbf{c}_j$  for some matrix  $M$ . One way to be more specific about this  $M$  is the following: Let  $P$  be the permutation matrix such that  $P\mathbf{Z} = (\mathbf{Z}_{\overline{\mathcal{C}}}, \mathbf{Z}_{\mathcal{C}})^{\top}$ . Then we have

$$\mathbb{E}(P\mathbf{Z} \mid \mathbf{Z}_{\mathcal{C}} = \mathbf{c}_j) = \begin{pmatrix} \Sigma_{\overline{\mathcal{C}}\mathcal{C}} \Sigma_{\mathcal{C}\mathcal{C}}^{-1} \\ I_{|\mathcal{C}|} \end{pmatrix} \mathbf{c}_j \implies M = P^{-1} \begin{pmatrix} \Sigma_{\overline{\mathcal{C}}\mathcal{C}} \Sigma_{\mathcal{C}\mathcal{C}}^{-1} \\ I_{|\mathcal{C}|} \end{pmatrix}$$

Hence, after conditioning, the expected value of the noise is

$$\mathbb{E}(\mathcal{E} \mid \mathbf{Z}_{\mathcal{C}} = \mathbf{c}_j) = (I - B) \mathbb{E}(\mathbf{Z} \mid \mathbf{Z}_{\mathcal{C}} = \mathbf{c}_j) = (I - B) M \mathbf{c}_j$$

Therefore, we can define

$$W = (I - B) M = (I - B) P^{-1} \begin{pmatrix} \Sigma_{\overline{\mathcal{C}}\mathcal{C}} \Sigma_{\mathcal{C}\mathcal{C}}^{-1} \\ I_{|\mathcal{C}|} \end{pmatrix} \in \mathbb{R}^{q \times |\mathcal{C}|}$$

Note that  $W$  depends on the nodes we are conditioning on, but not on their values. The rest of the proof is completely analogous to the previous one. ■

**Remark 4** *Although it is common to work with  $\mu_{\varepsilon} = \mathbf{0}$ , there is no need to assume that. In general, again using the Schur complement, we would have (let  $\mu = \mu_{\varepsilon}$ )*

$$\begin{aligned} \mathbb{E}(\mathbf{Z}_{\overline{\mathcal{C}}} \mid \mathbf{Z}_{\mathcal{C}} = \mathbf{c}_j) &= \Sigma_{\overline{\mathcal{C}}\mathcal{C}} \Sigma_{\mathcal{C}\mathcal{C}}^{-1} (\mathbf{c}_j - \mu_{\mathcal{C}}) + \mu_{\overline{\mathcal{C}}} = \Sigma_{\overline{\mathcal{C}}\mathcal{C}} \Sigma_{\mathcal{C}\mathcal{C}}^{-1} \mathbf{c}_j - \Sigma_{\overline{\mathcal{C}}\mathcal{C}} \Sigma_{\mathcal{C}\mathcal{C}}^{-1} \mu_{\mathcal{C}} + \mu_{\overline{\mathcal{C}}} \\ &= (\mu_{\overline{\mathcal{C}}} - \Sigma_{\overline{\mathcal{C}}\mathcal{C}} \Sigma_{\mathcal{C}\mathcal{C}}^{-1} \mu_{\mathcal{C}} \mid \Sigma_{\overline{\mathcal{C}}\mathcal{C}} \Sigma_{\mathcal{C}\mathcal{C}}^{-1}) \text{aug}(\mathbf{c}_j) \end{aligned}$$

So, if we define  $\mathbf{c}'_j = \text{aug}(\mathbf{c}_j) \in \mathbb{R}^{1+|\mathcal{C}|}$  and  $M$  accordingly, all the proof works analogously.



### Appendix C. Best approach for estimating a one-way fixed effect

In Section 4, we claimed that the best approximation of a one-way fixed effect is given by the mean-over-contexts estimator. Indeed, consider that following problem

$$\min_{\Gamma \in \mathbb{R}^m} f(\Gamma) := \|\mathbf{Y} - \Gamma \mathbf{1}^\top\|_F^2 \quad (4)$$

where  $\mathbf{Y} \in \mathbb{R}^{m \times n}$ ,  $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^n$ , and  $\|\cdot\|_F$  denotes the Frobenius norm of a matrix. By the linearity of trace we have

$$f(\Gamma) = \text{tr}(\mathbf{Y}^\top \mathbf{Y}) - 2 \text{tr}(\mathbf{1} \Gamma^\top \mathbf{Y}) + \text{tr}(\mathbf{1} \Gamma^\top \Gamma \mathbf{1})$$

Using the cyclic property of the trace and taking the gradient of  $f$  we get

$$\nabla f(\Gamma) = -2\mathbf{Y}\mathbf{1} + 2n\Gamma = \mathbf{0} \iff \Gamma = \frac{1}{n}\mathbf{Y}\mathbf{1}$$

That is,  $\Gamma_i = \frac{1}{n} \sum_j Y_{ij}$ , which corresponds to the mean-over-contexts estimator. Since  $f$  is strictly convex, this is the solution to Problem 4.

### Appendix D. SI is non-symmetric in the tensor case

In Section 4, we argued that SI is non-symmetric in the tensor case. The following example shows that this is true. In particular, completing within columns recovers the desired outcome while completing within rows does not. The key idea is that we have to flatten the tensor in one direction or another, and depending on the direction, linear span inclusion does or does not hold. Consider the following 3-order tensor

$$\mathbf{Y} = \left[ \begin{array}{cccc} Y_{111} & Y_{121} & Y_{131} & Y_{141} \\ Y_{211} & Y_{221} & Y_{231} & Y_{241} \\ Y_{311} & Y_{321} & Y_{331} & Y_{341} \\ Y_{411} & Y_{421} & Y_{431} & Y_{441} \end{array} \middle\| \begin{array}{cccc} Y_{112} & Y_{122} & Y_{132} & Y_{142} \\ Y_{212} & Y_{222} & Y_{232} & Y_{242} \\ Y_{312} & Y_{322} & Y_{332} & Y_{342} \\ Y_{412} & Y_{422} & Y_{432} & Y_{442} \end{array} \right] = \left[ \begin{array}{cccc} 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{array} \middle\| \begin{array}{cccc} 0 & 2 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{array} \right]$$

For the  $(i,j)$ -th entry, we use the notation  $\mathbf{Y}_{ij} = (Y_{ij1}, Y_{ij2}) \in \mathbb{R}^2$ . Suppose we have observed all the matrix except for the  $(4,4)$ -th entry. By regressing along rows, we are predicting the (flattened) first three rows of the 4th column  $(1, 0, 1, 1, 1)$  from the (flattened) first three rows of the 1st, 2nd and 3rd columns  $(1, 1, 1, 0, 1, 1)$ , etc. Thus, we obtain the following system of equations:

$$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 2 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix} \beta = \begin{pmatrix} 1 \\ 0 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$

This linear system is independent and we obtain  $\hat{\beta} = (1, 1, -1)$ . Thus, our prediction for the  $(4,4)$ -th entry is

$$\hat{\mathbf{Y}}_{44} = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \end{pmatrix} \hat{\beta} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \mathbf{Y}_{44}$$

Thus, regressing along rows works successfully. However, by regressing along columns we would obtain the following:

$$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \\ 2 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix} \alpha = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$

which is an inconsistent system. The least-squares solution is  $\hat{\alpha} = (0, 0.6, 0)$ , so the prediction would be

$$\hat{\mathbf{Y}}_{44} = \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix} \hat{\alpha} = \begin{pmatrix} 0 \\ 0.6 \end{pmatrix} \neq \mathbf{Y}_{44}$$

So this shows how SI works only in one direction. Moreover, this example can be related to our causal setup. Let  $Z_1 \rightarrow Z_2$  follow the SCM

$$\begin{cases} Z_1 = \varepsilon_1 \\ Z_2 = Z_1 + \varepsilon_2 \end{cases}$$

where  $\varepsilon_1, \varepsilon_2 \sim \mathcal{N}(0, 1)$  are independent noise variables. Suppose that we are observing both variables. Each column of  $\mathbf{Y}$  corresponds to conditioning on  $\mathbf{Z} = (c_1, c_2)$ . In particular, we have  $\mathbb{E}[(\varepsilon_1, \varepsilon_2) \mid (Z_1, Z_2) = (c_1, c_2)] = (c_1, c_2 - c_1)$ . The rows of  $\mathbf{Y}$  correspond to an intervention  $\text{do}(Z_1 = a_1, Z_2 = a_2)$ . The values of  $(c_1, c_2)$  and  $(a_1, a_2)$  are defined in Table 2, where "-" denotes

column ( $I_C$ )	$c_1$	$c_2$	row ( $I_A$ )	$a_1$	$a_2$
$j = 1$	1	0	$i = 1$	1	-
$j = 2$	0	1	$i = 2$	-	1
$j = 3$	1	1	$i = 3$	1	1
$j = 4$	0	0	$i = 4$	0	1

Table 2: Values used for defining  $i$ -th row and  $j$ -th column of  $\mathbf{Y}$ .

that we are not intervening on that variable. It is easy to check that  $\mathbb{E}(\mathbf{Z}(i) \mid I_C = j) = \mathbf{Y}_{ij}$ . Therefore, our matrix comes from a linear causal model. To see why SI only worked on one direction, recall that Equation (3) tells us that we can write

$$\mathbf{Y}_{ij} = P_i \mathbf{q}_j$$

for some  $P_i \in \mathbb{R}^{2 \times 3}$  and  $\mathbf{q}_j \in \mathbb{R}^3$ . Following the Proof of Theorem 1, we can construct these factors:

$$\begin{aligned} P_1 &= \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 1 \end{pmatrix} & P_2 &= \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} & P_3 &= \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} & P_4 &= \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \\ \mathbf{q}_1 &= (1 \quad 1 \quad -1) & \mathbf{q}_2 &= (1 \quad 0 \quad 1) & \mathbf{q}_3 &= (1 \quad 1 \quad 0) & \mathbf{q}_4 &= (1 \quad 0 \quad 0) \end{aligned}$$

Note that  $\mathbf{q}_4 \in \text{span}(\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3)$ . In fact,  $\mathbf{q}_4 = \hat{\beta}_1 \mathbf{q}_1 + \hat{\beta}_2 \mathbf{q}_2 + \hat{\beta}_3 \mathbf{q}_3$ , so the linear span inclusion is satisfied along columns. Nevertheless,  $P_4 \notin \text{span}(P_1, P_2, P_3)$ , so the linear span inclusion is not satisfied along rows.

## Appendix E. Polynomial case

In Section 4, we discussed the possibility of extending our results to non-linear SCMs. Here, we show an example where the latent variables follow a linear SCM and there is a polynomial link function mapping the latent variables to the observed ones. For simplicity in notation we use  $X = \mathbf{Z}_X$ ,  $A = \mathbf{Z}_A$ , and  $C = \mathbf{Z}_C$ . Note that  $\mathcal{A}$  and  $\mathcal{C}$  are the same for all rows and columns, respectively.

**Example 2** Consider the following SCM:

$$\begin{cases} A = \varepsilon_A \\ C = B_{AC}A + \varepsilon_C \\ X = f(A, C) + \varepsilon_X, \quad f \in \mathbb{R}_d[Y_1, Y_2] \end{cases}$$

Since there is linear relation between  $A$  and  $C$ , we can still use the Schur Complement. For a gaussian variable  $Z \sim \mathcal{N}(\mu, \sigma^2)$ , we have  $\mathbb{E}(Z^k) = p_k(\mu, \sigma)$  for some polynomial  $p_k \in \mathbb{R}_k[T_1, T_2]$ . Therefore,  $\mathbb{E}(X(a_i)|C = c_j)$  is a polynomial in  $a_i$  and  $c_j$ . In particular, we can write it as

$$\begin{aligned} \mathbb{E}(X(a_i)|C = c_j) &= q_0(c_j) + q_1(c_j)a_i + \dots + q_{d-1}(c_j)a_i^{d-1} + q_d(c_j)a_i^d \\ &= \begin{pmatrix} 1 & a_i & \dots & a_i^d \end{pmatrix} \begin{pmatrix} q_0(c_j) \\ q_1(c_j) \\ \vdots \\ q_d(c_j) \end{pmatrix} \end{aligned}$$

for some polynomials  $q_k$ .<sup>5</sup> Therefore, for a polynomial of degree  $d$  we have a factor model of degree  $d + 1$ , so we could also use Synthetic Interventions in this case.

If we had multiple fixed action nodes  $A_1, \dots, A_r$  and a polynomial of degree  $d$  from the latent nodes to  $X$ , we could use the same idea and express  $\mathbb{E}(X(a_i) | C = c_j)$  as a polynomial on  $a_i^1, \dots, a_i^r$ . The number of monomials in  $r$  variables of degree no greater than  $d$  is  $\binom{r+d}{d}$ , so that would be the rank of our factor model. However, it is not clear how our decomposition would look like if allowed intervening on different nodes, because the  $q_k$ 's would depend on these indices. We leave this question open as a future direction.

### E.1. Going beyond polynomials

**Example 3** Consider now the following SCM:

$$\begin{cases} A = \varepsilon_A \\ C = A + \varepsilon_C \\ X = \exp(AC) \end{cases}$$

After conditioning and intervening, we get that

$$X(a)|C = c \sim \text{Lognormal} \left( a \left( a + \frac{c}{2} \right), \frac{a^2}{2} \right)$$

5. In this case, we could defined this polynomial more explicitly as a function of  $\mathbb{E}(\varepsilon_c^k | C = c_j)$ , but this is not necessary for expressing the factor model.

Hence,

$$\mathbb{E}(X(a)|C = c) = \exp\left(a\left(a + \frac{c}{2}\right) + \frac{a^2}{4}\right) = \exp\left(\frac{5}{4}a^2\right) \exp\left(\frac{ac}{2}\right)$$

But this expression cannot be expressed as a low-rank factor model.

Example 3 suggests that it might be difficult to obtain low-rank factor models for SCMs that involve functions "more complicated" than polynomials. Moreover, Example 2 shows that the rank of the factor model increases as the degree of the polynomial increases.

Note that the expression obtained is the exponential of a polynomial in  $a, c$ . If we applied an entry-wise logarithm to our matrix, we could apply the SI estimator and then take the exponential of the outcome observed. With this idea in mind, we could think about developing SI theory for feature spaces. However, this goes beyond the scope of this work and we leave it for future discussion.

## Appendix F. Hypothesis tests

In Section 4, we discussed the necessity to test the validity of our assumptions. Here, we show how our causal modeling may also be useful for constructing hypothesis tests. For simplicity, assume that we have observed  $n_s$  independent samples for each observed entry, i.e. for each  $(i, j) \in \Omega$ . Assume that  $Y_{ij} \sim \mathcal{N}(\mu_{ij}, \sigma_{ij}^2)$  and let  $\bar{Y}_{ij} \sim \mathcal{N}(\mu_{i,j}, \frac{\sigma_{ij}^2}{n_s})$  denote the average of these samples. We can estimate the sample variance as follows  $S_{ij}^2 = \frac{1}{n_s-1} \sum_{k=1}^{n_s} (Y_{ij}^{(k)} - \bar{Y}_{ij})^2$ .

### F.1. Homoscedasticity within rows/columns

Before discussing the different decompositions obtained in Section 4, is it worth noting that the SCMs we studied give rise to certain structure in the variance of our matrix. In particular, Prop. 2 leads to homoscedasticity within the entries of the same row ( $\sigma_{i,j} = \sigma_{i,j'}$ ). Indeed, looking at the Schur Complement, the variance of the noise after conditioning on some variables depends on the variables we condition on, but not on the actual values they take. On the other hand, in Corollary 3 we have the same SCMs after the interventions for different rows, which implies homoscedasticity within columns ( $\sigma_{i,j} = \sigma_{i',j}$ ).

This lead to a classical test for comparing the variances of two independent samples. For example, for testing homoscedasticity within column  $j$  we consider

$$\begin{cases} H_0 : \frac{\sigma_{i,j}^2}{\sigma_{i-1,j}^2} = 1 & \forall i \geq 2, j \\ H_1 : \frac{\sigma_{i,j}^2}{\sigma_{i-1,j}^2} \neq 1, & \exists i, j \end{cases}$$

So we can use the following estimator

$$F = \frac{S_{i,j}^2}{S_{i-1,j}^2}$$

and we have  $F | H_0 \sim F_{n_s-1, n_s-1}$ . A test for homoscedasticity within row  $i$  would be analogous.

## F.2. Test for fixed effects

Suppose that we are in the case of Corollary 3. How can we test if FE is appropriate for our matrix? For simplicity in notation, suppose that we have observed all the entries of  $\mathbf{Y} \in \mathbb{R}^{m \times n}$ . For a partially observed entries, we would use these tests for the observed entries. With this notation, all the matrices with a two-way fixed effect are defined by the following equations.

$$Y_{ij} + Y_{i'j'} - Y_{i'j} - Y_{ij'} = 0 \quad \forall i \neq i', j \neq j' \quad (5)$$

Many of this equations are linearly dependent. In fact, all of them can be expressed as a linear combination of equations

$$E_{ij} : Y_{11} + Y_{ij} - Y_{i1} - Y_{1j} = 0, \quad i = 2, \dots, m, j = 2, \dots, n$$

Indeed, by computing  $E_{ij} + E_{i'j'} - E_{i'j} - E_{ij'}$  we get Equation (5). So we only need to check  $(m-1)(n-1)$  equations instead of  $m(m-1)n(n-1)$ . However, one drawback about using Equation (F.2) for constructing hypothesis tests is that they heavily rely on the (1,1)-th entry. So we can consider a more robust set of equations:

$$\tilde{E}_{ij} : Y_{i-1,j-1} + Y_{i,j} - Y_{i,j-1} - Y_{i-1,j} = 0, \quad i = 2, \dots, m, j = 2, \dots, n$$

Let  $\mu_{i,j} = \mathbb{E}(Y_{i,j})$ , the hypothesis test for FE is the following.

$$\begin{cases} H_0 : \mu_{i-1,j-1} + \mu_{i,j} - \mu_{i,j-1} - \mu_{i-1,j} = 0, & \forall i, j \\ H_1 : \mu_{i-1,j-1} + \mu_{i,j} - \mu_{i,j-1} - \mu_{i-1,j} \neq 0, & \exists i, j \end{cases}$$

Therefore, we can consider the following statistic:

$$\hat{T}_{ij} = \frac{\bar{Y}_{i-1,i-1} + \bar{Y}_{i,j} - \bar{Y}_{i,j-1} - \bar{Y}_{i-1,j}}{\sqrt{\frac{1}{n_s} \left( S_{i-1,j-1}^2 + S_{i,j}^2 + S_{i,j-1}^2 + S_{i-1,j}^2 \right)}}$$

The key idea is that, in Corollary 3, we have homoscedasticity within rows. This makes our test similar to a Welch's test, and we can approximate the distribution of  $\hat{T}_{i,j} \mid H_0$  as  $t_{\hat{\nu}}$  where

$$\hat{\nu} = (n_s - 1) \frac{\left( (S_{i,j}^2 + S_{i-1,j}^2) + (S_{i,j-1}^2 + S_{i-1,j-1}^2) \right)^2}{(S_{i,j}^2 + S_{i-1,j}^2)^2 + (S_{i,j-1}^2 + S_{i-1,j-1}^2)^2}$$

**Multiple comparisons problem.** We have  $N = (n-1)(m-1)$  different null hypothesis, corresponding to each  $(i, j)$  pair for  $i, j \geq 2$ . If we want to achieve a familywise error rate (FWER) of  $\alpha$ , we can reject a null hypothesis if we get a p-value lower than  $\frac{\alpha}{N}$  (Bonferroni correction), or  $1 - (1 - \alpha)^N$  (Šidák correction).

**Empirical Results.** In Fig. 7 we show the how this estimator is useful for testing a fixed-effect behaviour. In particular, for each plot, we simulate 200 matrices satisfying the assumptions from Corollary 3 and 200 matrices satisfying the assumptions from Prop. 2. More specifically, we use the `causalDag` Python library to generate a DAG with random weights (we use the default parameters). We fix the node we observe ( $\mathbf{Z}_{\mathcal{X}}$ ) and the context variables  $\mathbf{Z}_{\mathcal{C}}$ . For each simulated matrix, we

condition on a different set of values ( $\mathbf{z}_c$ ), sampled from a uniform distribution over  $[0, 10]$ . The nodes where we intervene on change for each matrix. We perform a  $\text{do}$ -intervention on each node with probability 0.2 and the values we used are also sampled from a uniform distribution over  $[0, 10]$ . For the matrices where the fixed effect behavior is not satisfied, we do not intervene on the same nodes for all the matrix, so we follow the procedure mentioned before for each row.

We observe how, as the size of the matrix and the number of observed samples increase, the statistic we propose can fully distinguish the two types of matrices.

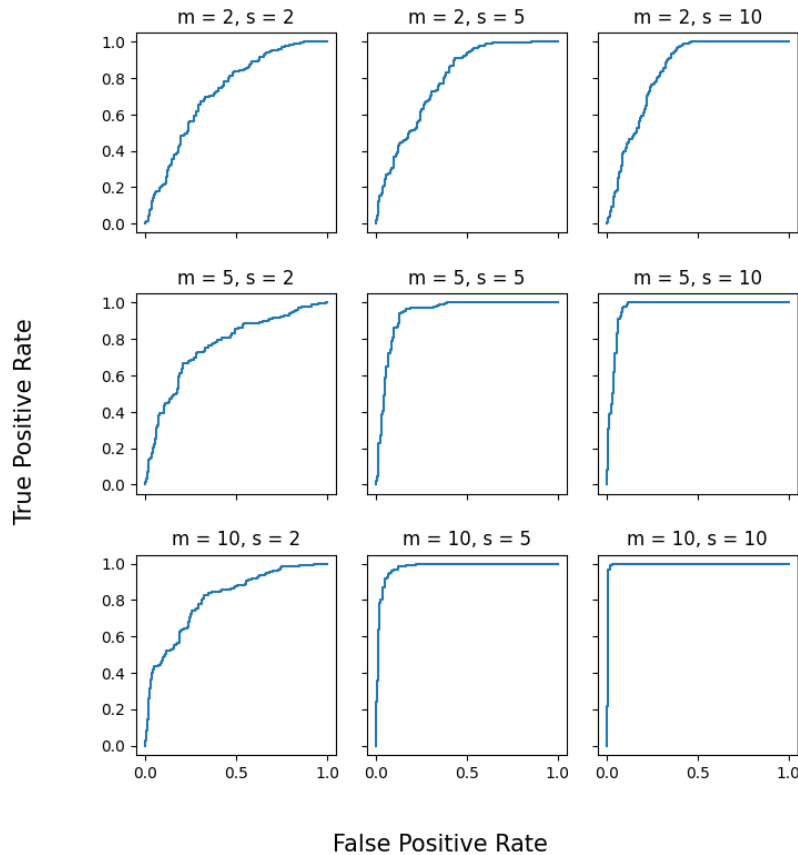


Figure 7: **ROC curves.**  $m$  denotes the size of the matrix, i.e.  $\mathbf{Y} \in \mathbb{R}^{m \times m}$ .  $s$  denotes the number of samples observed per entry.

### E.3. Test for mean-over-contexts

To be precise, before testing whether or not we our matrix follows a two-way fixed effect model, we should test whether it has only a one-way fixed effect. That is, test if we can use mean over actions or contexts. For example, to test whether we only have a fixed effect along rows, instead of Equation 5 we would have

$$Y_{ij} - Y_{ij'} = 0 \quad \forall i, j \neq j'$$



And we can repeat all the reasoning from the previous section to obtain an appropriate estimator. The case of a fixed effect along columns would be completely analogous.

#### **F.4. Test for Synthetic Interventions**

[Agarwal et al. \(2020\)](#) propose a hypothesis test for the SI estimator, so we could use it for our particular case to check the linear span inclusion assumption. Moreover, as mentioned before, assuming Gaussianity as in [Prop. 2](#) we can also test homoscedasticity within rows. Therefore, constructing such tests is an interesting future direction.

## Appendix G. Further analysis on PRISM Repurposing dataset

### G.1. SVD analysis

In Section 5, we argued that getting high  $R^2$  was difficult because the PRISM data is relatively noisy. Here we analyze the Singular Value Decomposition (SVD) of our matrix  $\mathbf{Y}$  shown Figure 4. It is important to mention that here we are showing the SVD computed on all the matrix. That is, we are also using the entries that were supposed to be missing and that we want to impute as a matrix completion task. Therefore, the results shown in this section are merely an analysis of our data, but we are not going to use them for the imputation task.

In Figure 8, we show the Singular Values of  $\mathbf{Y}$ . We can see that they are relatively diffuse, suggesting that there is a considerably amount of noise in the data. For example, the explained variance by the top 10 SVD vectors  $\left(\frac{\sum_{i \leq 10} \sigma_i^2}{\sum_i \sigma_i^2}\right)$  is just 78%.

The Eckart-Young theorem (Eckart and Young, 1936) states that the best rank  $r$  approximation to a matrix is given by truncating its SVD to the top  $r$  singular values. We can use this result to establish upper bounds in performance that we should expect from matrix completion approaches. For instance, mean-over-contexts produces a rank 1 approximation of our matrix, so this approximation must be worse (i.e. smaller  $R^2$ ) than the rank 1 truncated SVD.

In Figure 8, we also show the  $R^2$  values for the truncated SVD using the missing data pattern like the one from Figure 5, with 25% of missing entries. It is important to remark that in this case we are computing the SVD using all the matrix  $\mathbf{Y}$  (including missing entries), so this is not a valid completion approach. It should be understood as the optimal performance we can expect.

Generally, it is helpful to see how the first singular values pairs look like. Let  $\mathbf{S}_1$  be rank 1 truncated SVD of  $\mathbf{Y}$ . In Fig. 9 we observe that  $\mathbf{S}_1$  has almost constant rows, so it seems to be capturing the fixed effect in that direction. This indicates that the fixed effect along rows is much more important than the fixed effect along columns. Therefore, we may expect that the mean-over-contexts is going to be a considerably good baseline. We also see how in  $\mathbf{Y} - \mathbf{S}_1$  we lose a big part of this fixed effect along rows, and our matrix seems to have less variability. This suggests that subtracting the mean-over-contexts may lead us to better predictions (see Section 4). Finally, SVD2 also seems to capture some of this fixed effect behavior, but in a smaller order of magnitude.




Figure 8: **Difficulty in getting high  $R^2$  values.** The singular values of  $\mathbf{Y}$  present a relatively heavy tail. This makes it difficult to obtain good metrics in our predictions. The truncated SVD of rank  $r$  is the best rank  $r$  approximation of a matrix. Therefore, the  $R^2$  obtained by the truncated SVD may be considered as an upper bound for our completion algorithms. We should not expect to achieve an  $R^2$  of 0.90 as this would involve finding the best approximation of rank 80.

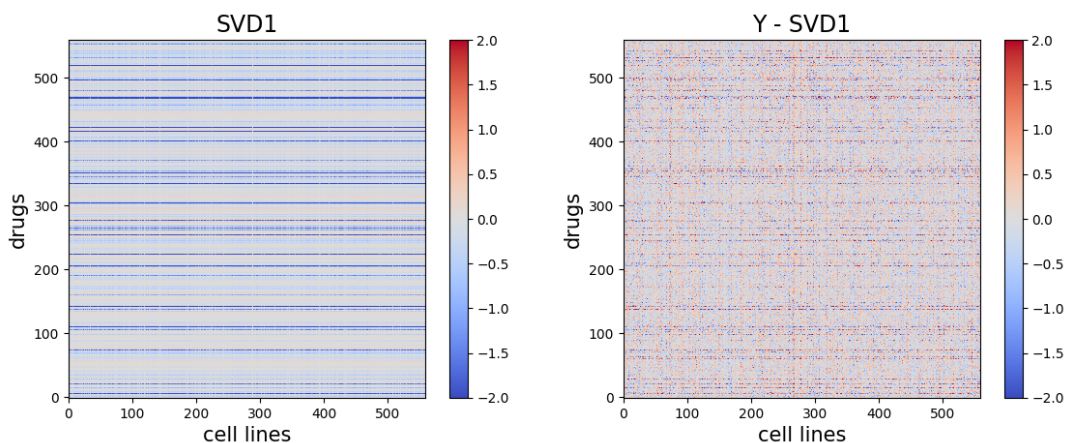


Figure 9: **One-way fixed effect behavior.** The truncated SVD of rank 1 ( $\text{SVD1} = \mathbf{S}_1$ ) has approximately constant rows. This suggests that there is a fixed effect for actions so mean-over-contexts should be a relatively good estimator.

## G.2. Missing pattern in CMAP dataset

As we mentioned in Section 5.1, in biology applications we usually do not have a missing at random pattern. In Fig. 10 we show the entries observed for the CMAP dataset. We use this example as a starting point for the missing data pattern considered in Section 5.1.




Figure 10: **Missing pattern NOT at random.** Availability matrix for CMAP dataset. A black rectangle means that we have observed that entry. Columns have been sorted from left to right according to the number of available entries (for each column). In a similar way, rows have been sorted from bottom to top. Figure obtained from [Squires et al. \(2022\)](#).

### G.3. An alternative missing data pattern

The missing data patterns from Section 5.1 may seem very artificial, as we have a constant number of observed rows and columns for each missing entry. In Fig. 11 we propose an alternative pattern, still following the motivation from the CMAP dataset (see Appendix G.2), and we show the performance of different algorithms for this setting. The results are not significantly different to the ones from Fig. 6.




Figure 11: **Non-square missing data pattern.** Here, we do not have a constant value of rows and columns for each missing entry.

#### G.4. Performance of the algorithms depending on the amount of data observed

In Section 5.2, we showed the performance of different matrix completion approaches for a specific number of observed entries. The experiment we show here consists on varying the number of rows and columns observed per missing entry. This parameter will range from 5 until 275. For each case, the square corresponding to the missing values has constant size  $284 \times 284$ . The reason for having constant size square (same number of missing entries) is that the MSE is computed on the same number of samples.

We have used the following bootstrap technique. For each case we shuffle the rows and columns of the full matrix and take a sub-matrix of our desired size (e.g.  $289 \times 289$ ). Then, we take the corresponding missing data pattern and compute the  $R^2$  for each algorithm. We repeat the same process for 20 different shuffles. See Fig. 12. As a takeaway, it is worth noting that Synthetic Interventions benefits from subtracting the mean-over-contexts, especially in the low data regime.

Note that we can think about Fig. 6 as taking slices from Fig. 12. In particular, Fig. 6 shows that Collaborative Filtering approach does not improve the baseline algorithms and Kernel Linear Regression is not significantly better than the standard Synthetic Interventions approach. For that reason, we do not include them in Fig. 12.

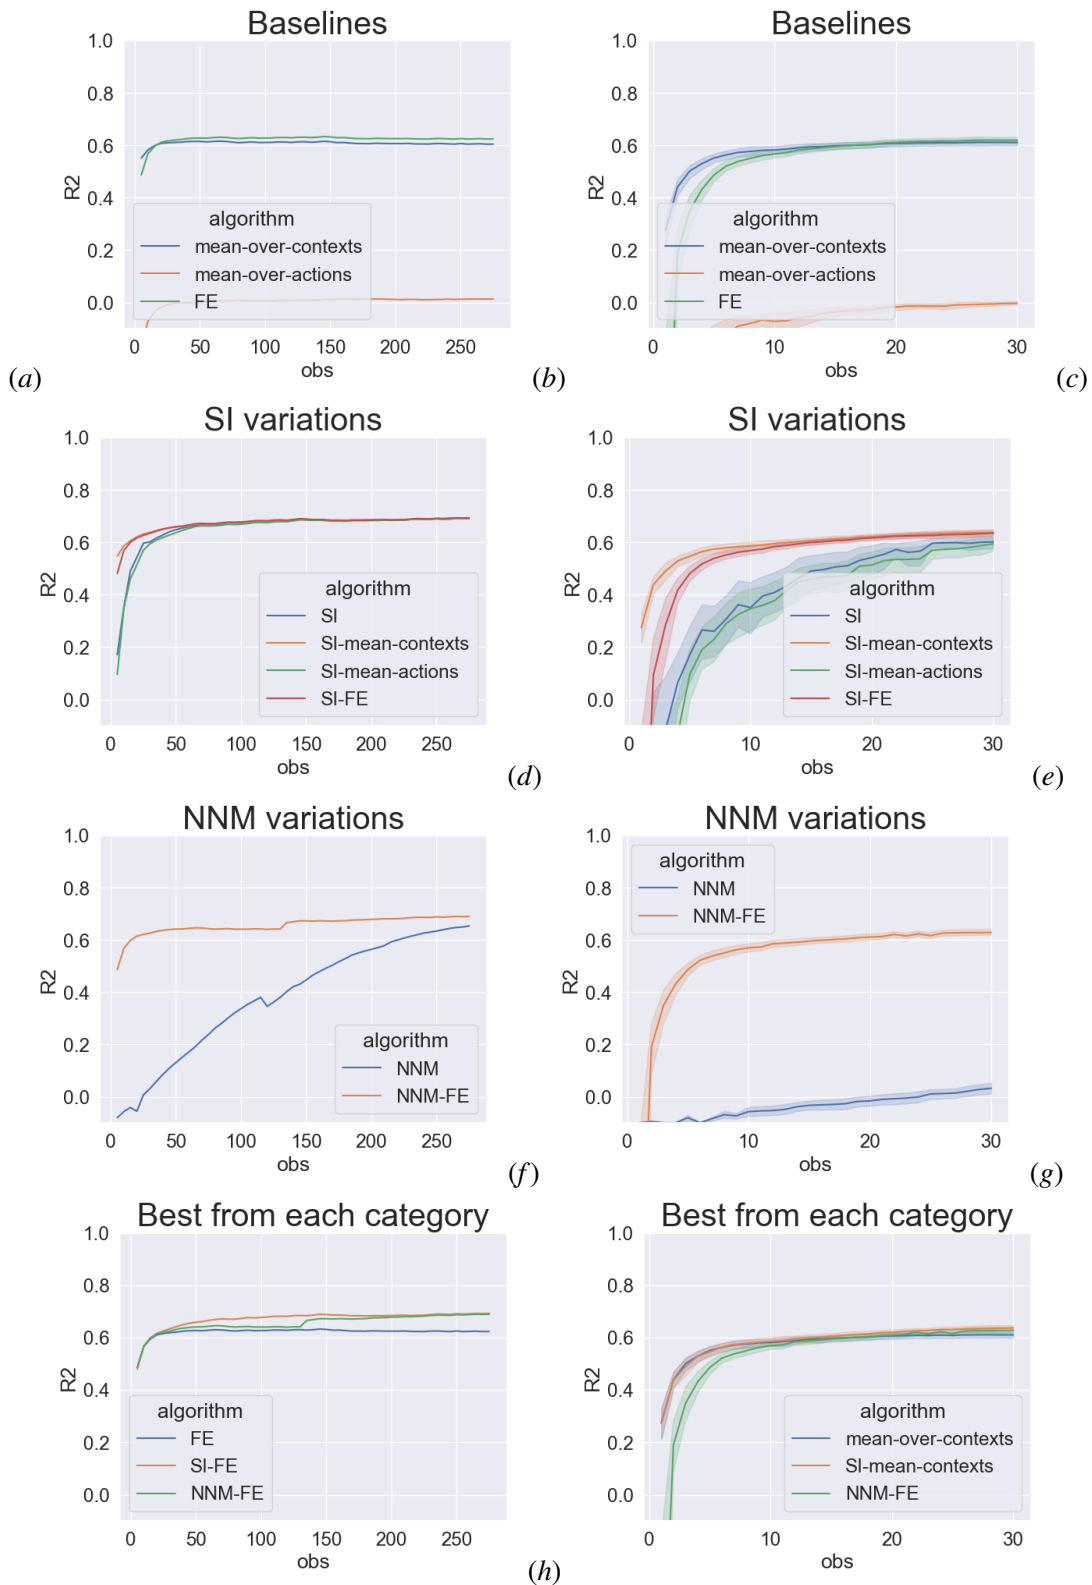


Figure 12: **Dependence on number of observations.** "obs" denotes the number of rows and columns observed for each missing entry. Figures from the right-hand side are focused on the low-data regime. We can see some spikes in the curves corresponding to NNM, these are produced by a change in the regularization parameter (e.g. from  $10^{-3}$  to  $10^{-4}$ ).

### G.5. Performance of all the matrix completion algorithms considered

In Section 5.2, we showed the performance of some matrix completion approaches. Here, in Fig. 13, we include more approaches and some variations of the previous ones.

In particular, we see that the nearest neighbors approaches are not significantly better than the baselines. Doubly Robust NN is the method proposed in Dwivedi et al. (2022). CF10 stands for a 10-nearest neighbor Collaborative Filtering approach. Finally, it is worth noting that SI and its variations outperform the other approaches.




Figure 13: **All the matrix completion approaches considered.** Here we show the performance of all the algorithms considered, using the missing pattern from Figure 5.

### G.6. Removal of "killer drugs"

In this section, we run the same experiments on certain portions of the whole matrix  $\mathbf{Y}$ . Looking back at our original matrix in Fig. 4, we observe that there are some blue (negative) rows that seem to be constant along all the columns. This would mean that this particular drug is killing all the cell lines.

In light of this effect, we call a drug a "killer drug" if it kills more than a certain percentage of cell lines. It is not clear which threshold we should use. In Fig. 17, we show the number of "killer drugs" depending on the threshold we define.

For the following results, we used a threshold of 80%. The resulting matrix can be seen in Fig. 14. Note that we have removed almost all the blue lines that we mentioned before. Now the fixed effect along rows does not seem that clear. In particular, the MSE of the baseline decreases significantly compared to the previous setup. This means that the task of completing the task is *easier* for the baseline model. In consequence, getting high  $R^2$  values is going to be much harder.



In Fig. 15 we repeat the same experiment from Fig. 8 for this new data. As expected, now the  $R^2$  value we obtain are much lower. Finally, we show the performance of the different matrix completion approaches on this matrix in Fig. 16.

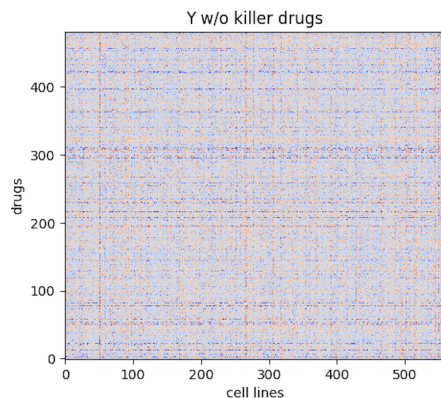


Figure 14: **Matrix obtained when we remove "killer drugs"**. Here, a "killer drug" is a drug that kills more than 80% of cell lines.




Figure 15: **Getting high R2 values is even more difficult than before** (compare it with Fig. 8). After removing the "killer drugs", MSE(average) decreases from 0.87 to 0.28. This means that the completion task is easier for this baseline, so it is more difficult to obtain high  $R^2$  values.

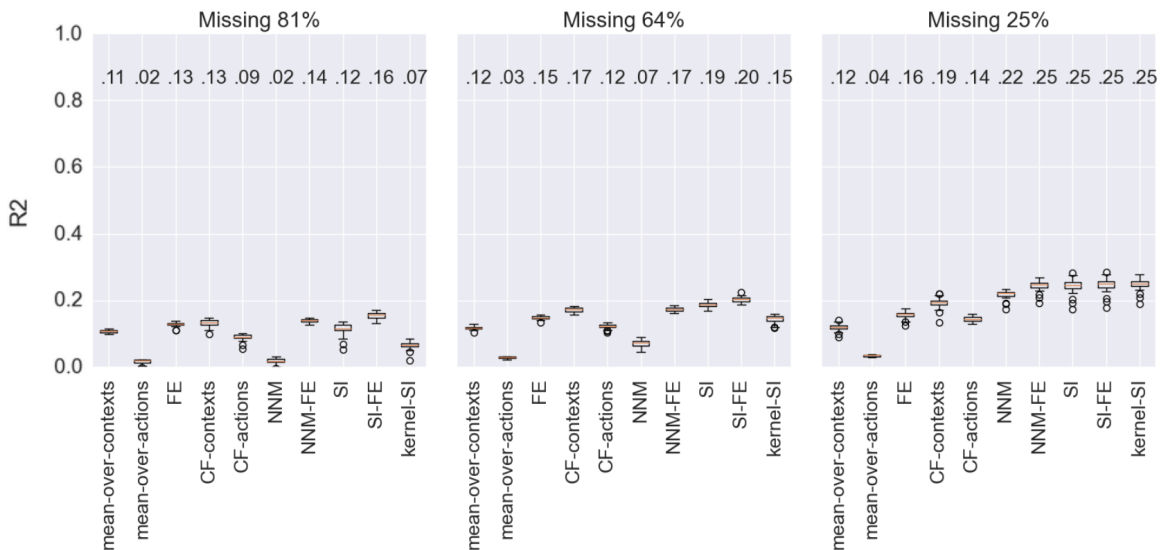


Figure 16: **Performance after removing "killer drugs"**. The gap between mean-over-actions and SI is similar from the one in Fig. 6, despite having more room for improvement.

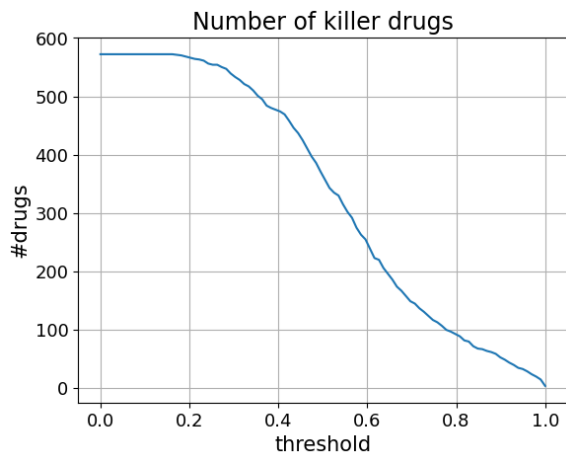


Figure 17: **Definition of a "killer drug"**. Here we show the number of drugs that kill more than *threshold* % of cell lines. We consider that a cell line is dead if the viability score is negative. In our experiments, we used a threshold of 80%.