
Compositional Abstraction Error and a Category of Causal Models

Eigil F. Rischel¹

Sebastian Weichwald²

¹Department of Computer & Information Sciences, University of Strathclyde, United Kingdom

²Department of Mathematical Sciences, University of Copenhagen, Denmark

Abstract

Interventional causal models describe several joint distributions over some variables used to describe a system, one for each intervention setting. They provide a formal recipe for how to move between the different joint distributions and make predictions about the variables upon intervening on the system. Yet, it is difficult to formalise how we may change the underlying variables used to describe the system, say moving from fine-grained to coarse-grained variables. Here, we argue that compositionality is a desideratum for such model transformations and the associated errors: When abstracting a reference model M iteratively, first obtaining M' and then further simplifying that to obtain M'' , we expect the composite transformation from M to M'' to exist and its error to be bounded by the errors incurred by each individual transformation step. Category theory, the study of mathematical objects via compositional transformations between them, offers a natural language to develop our framework for model transformations and abstractions. We introduce a category of finite interventional causal models and, leveraging theory of enriched categories, prove the desired compositionality properties for our framework.

1 INTRODUCTION

With a causal model we aim to predict future observations of a system under the same conditions that held true when we devised the model (observational distribution), or when the system is subjected to external manipulating forces (interventional distributions), or infer what the observations would have been like had the context been different (counterfactuals). Following the common aphorism “all models are wrong”, we do not and cannot regard any causal model

as a precise description of some system but instead as an approximate description that is pragmatically useful to us. Besides the question of “usefulness”, which inevitably depends on modelling goals, there is another one: How well does our model approximate reality?

The ground-truth may be unattainable. Therefore, an exact quantification and characterisation of how well a model describes reality may be beyond reach. It is possible, however, to assess a model relative to another model. Recent approaches aim to formalise this by considering the following question about transformations that link two models [Rubenstein et al., 2017, Beckers and Halpern, 2019, Beckers et al., 2019]: How well does a transformed, simpler and higher-level causal model agree with or approximate another, more detailed and lower-level model? An answer to this question is pragmatically useful. It allows us to bound the error of a model used in practice relative to the corresponding most accurate state-of-the-art reference model available (which substitutes ground-truth). Based on the error relative to one reference model, we can make a principled decision between candidate causal models of varying aggregation level and degree of simplification.

Model development is an iterative process. For example, specialists may propose approximations and transformations to only parts of a complex causal model, which we incorporate in a joint model and which we may refine further as we gain new insights. It is desirable that an account of model transformation facilitate such modular step-by-step simplifications. Existing work partly addresses this: Rubenstein et al. [2017, Lemma 5] prove that the notion of exact transformations is transitive: if M can be exactly transformed into M' , and M' into M'' , then the transformation from M to M'' is also exact; Beckers and Halpern [2019, Theorem 3.9] prove an analogous result for a stricter notion of transformations; and Beckers et al. [2019, Section 5] discuss the problem of composing approximations and abstractions, instead of exact transformations, in different order.

Complementing this line of research, we argue that a key

desideratum for an account of model abstraction is that the measure of error be compositional. We introduce a framework based on a category of interventional causal models, as category theory is a natural language for the study of models and compositional transformations between them. Extending the result of Rubenstein et al. [2017] to in-exact transformations, we prove that our proposed error notion is compositional: when we approximately transform M into M' , and M' into M'' , then an approximate transformation M to M'' exists and its error admits a natural bound in terms of the error of the two component transformations.

The categorical perspective is instructive and establishes an interdisciplinary common ground between applied category theory and causal modelling. We prove the well-definedness of interventional distributions by string diagram surgery, thereby reproducing classical results in our category theoretical framework. In the following subsections, we discuss conceptual issues in causal modelling that motivate our framework and can be addressed and formalised within it. In Section 2, we introduce our compositional framework, discuss design choices and its desirable features. We discuss how a notion of implemented models may enable us to reasonably talk about bounding the error of a causal model relative to ground-truth and based on interventional experiments in Section 3. Section 4 concludes our article.

1.1 VARIABLES IN STATISTICAL MODELS

In probabilistic models, the state of a system is represented by random variables. The values are determined by some measuring procedure. For example, we may model the daily warmth situation in an enclosed space using the following variables: A_1 and A_2 denoting the settings of two air conditioners (ACs) as transcribed to our spreadsheet at 8 o'clock in the morning, and T denoting the reading of a room thermometer at noon. Via a model over (A_1, A_2, T) , we may predict the AC configuration (A_1, A_2) given $T = t$.

We often decide implicitly which variables to use to describe a system. This decision is constrained by measurement ability and accuracy, computational and storage constraints, and pragmatic considerations. Yet, the choice is not arbitrary and we do rely on our descriptors to capture relevant aspects of reality. For example, to determine the value of T we may choose to measure a proxy of the average velocity of gaseous particles in a room with a thermometer, reading it off at 8:01:32 every day, and rounding to 2 digits. By doing so, we implicitly decide against measuring all particles' momenta and including billions of variables, and we fix which world states are (to be) considered indistinguishable for our modelling purposes. To make a conscious informed decision on the measurement protocols and variables to use, we need to understand how models using different variables relate, which is the subject of this article.

1.2 INTERVENTIONS IN CAUSAL MODELS

Interventional causal models make predictions about A_1, A_2 not only for given/observed $T = t$ but also for situations when we manipulate the value of T . In our example, we may write $\text{do}(T = t)$ for an intervention indicating that the value t of T did not come about naturally by following the measurement procedure outlined above, but instead was set externally. The intervention $\text{do}(T = t)$ may correspond, for example, to a situation where we a) entered the room shortly before noon and lit a lighter next to the thermometer before its reading was obtained, or b) taped the needle so the thermometer would have a fixed reading, or c) changed for a digital thermometer with a zero-digit precision display.

How a causal model relates to the tangible world thus depends not only on the measurement procedures to obtain the variable values but also on the physical implementations of interventions. Both determine how our model variables, which serve as explanans in our causal description, relate to reality or a state-of-the-art reference model instead of the unattainable ground-truth. There are multiple measurement procedures to obtain the variables' values. Additionally, the physical procedure corresponding to an intervention $\text{do}(T = t)$ in our model is ambiguous. For our causal model transformations, we thus require the modeller to make the content of a causal model that is to be preserved explicit and to specify how observations and interventions map between the transformed model and the reference model.

1.3 AMBIGUOUS INTERVENTIONS

It is problematic and unrealistic if—as is customary—the predictions a model makes about the effects of interventions do not depend on how the intervention is being implemented [Spirtes and Scheines, 2004]. Even if interventional data were available, the definition of causal variables may be underdetermined [Eberhardt, 2016].

Therefore, we require the modeller to make explicit how high-level interventions are implemented on the low-level. Ambiguity in the low-level implementation of interventions is encoded in an intervention kernel. For example, we can encode if we only ever intervened on total cholesterol (TC) by prescribing a certain restricted set of diets with comparable effects on the level of low-density (LDL) and high-density lipoprotein (HDL). Whether we allow only a restricted set of diets or any diet carries over, via the intervention kernel, to a lower or higher abstraction error when transforming a model and replacing LDL and HDL by TC.

1.4 WHY COMPOSITIONALITY?

Without compositionality, we cannot assess a model's error by comparing against the predecessor model, but instead need to evaluate each model relative to the reference model.

The following example illustrates how this lack of modularity deceives us when iteratively simplifying and abstracting different parts of a model. To compare the interventional distributions implied by the transformed model to those of the reference model, we here use the KL-divergence as a common measure of how one distribution differs from a reference distribution. We motivate why an information theoretic measure should be used in Section 1.5.

Imagine a model M of two ACs affecting the temperature in a room. The variables are A_1, A_2 for the ACs settings, and T for the temperature. The causal structure is $A_1 \rightarrow T \leftarrow A_2$, and we define the mechanisms by stating the respective kernels. Assume both ACs have settings $\{1, \dots, 100\}$ and the temperature scale is $\{1, \dots, 100\}$. Suppose that

- $P(T = n \mid A_1 = 100, A_2 = 100) = Z/n^2$,
- $P(T = n \mid A_1 = 100, A_2 \neq 100) = Z'/n^3$,
- $P(T = n \mid A_1 \neq 100) = Z''/100^n$,
- the settings of the ACs are independently and uniformly distributed in the observational setting, and
- Z, Z', Z'' are normalising constants.

Consider an abstraction M' of M , where the influence of the AC A_2 is simplified away, that is, we remove the arrow $T \leftarrow A_2$. Define the simplified model such that the distributions, $P(T \mid A_1), P(A_1), P(A_2)$ agree with the reference model M as closely as possible. As the simplified model requires that T and A_2 be independent, the two models disagree on $P(T \mid A_2)$. We compare the two models' predicted distribution for $T \mid \text{do}(A_1 = 100, A_2 = 100)$. A calculation shows that the KL-divergence between the two models' interventional distributions is a fairly small 0.22.

Consider another abstraction M'' of M' where the influence of the AC A_1 is also simplified away. The KL-divergence between the predictions of M' and M'' on $T \mid \text{do}(A_1 = 100, A_2 = 100)$ is again a fairly small 0.39.

We may expect that abstracting M as M'' should be permissible, since both abstractions, from M to M' and from M' to M'' were permissible. While small errors may accumulate, we expect the error of the abstraction M to M'' to be bounded in terms of the errors of abstracting M as M' and M' as M'' . One way to formalise this expectation is to impose the triangle inequality such that the transformation error of M to M'' is bounded by $0.39 + 0.22 = 0.61$. The KL-divergence between the predictions of M and M'' , however, is more than twice that: 1.52. For $\text{do}(A_1 = 100, A_2 = 100)$, M predicts a temperature above 1 in 39% of cases, while M'' predicts a temperature above 1 only in 0.01% of the cases. The errors of the individual abstraction steps are not indicative of how well the abstracted model M'' approximates the reference model M . The discrepancy between the individual and overall abstraction error is unbounded; in Appendix A. we construct for any $\epsilon, K > 0$, a situation

where the individual abstractions incur KL-divergences $\leq \epsilon$, while the overall abstraction incurs a KL-divergence $> K$.

This is a severe hurdle to the development of causal models. We often do not have access to ground truth but only have a model we think is reasonably accurate and are considering replacing it with another, perhaps simpler, one. If the errors are not compositional, as in the above example, we cannot ensure that the abstracted model M'' closely approximates the reference model (or ground-truth) M by enforcing a small approximation error relative to a previously established good approximation M' . This limits how informative the efforts to develop M' and to empirically validate its close resemblance of M are about M'' (see also Section 3).

Therefore, compositionality is a desideratum for our notion of causal model abstraction. Our framework is compositional: the errors of the two abstractions above are 0.2 and 0.22, and the composite abstraction error is $0.37 \leq 0.42$.

While abstraction examples similar to those discussed by Rubenstein et al., Beckers and Halpern, Beckers et al. can be expressed within our framework, the above example – where the “simplification” consists of deleting causal arrows, rather than, for example, reducing the number of variables – is instructive to exhibit the failure of compositionality.

1.5 ASSESSING MODEL ABSTRACTION ERROR

Often, we cannot establish an exact correspondence between two modelling levels. The conditions for the transformations and abstractions discussed by Rubenstein et al. [2017] and Beckers and Halpern [2019], and for modelling equilibria of a time-evolving process [Janzing et al., 2018] are restrictive. Therefore, we also wish to characterise transformations in which the transformed model only approximates the causal relationships in the reference model. Still, we require the two models, where one can be viewed as a transformation of the other, to stand in a well-defined relationship. A natural idea is to ask for an approximate transformation that preserves the causal structure *up to some error*. This idea has been proposed, for instance, by Beckers et al. [2019].

In Beckers et al.'s approach, the measure of abstraction error ultimately depends on a choice of metric on the underlying set of outcomes. The advantage of such an approach is obvious: it allows us to adjust the metric to optimally capture those aspects of the model that are of interest.

We see two problems with this approach. First, choosing a metric on the set of high-level variables requires that we already have chosen that abstraction level. In order to assess an abstraction we are required to decide how important the high-level variables we have just invented are. This is an unnatural requirement for comparing and finding candidate abstractions. Second, having to choose a metric at all requires detailed knowledge about the system and the

model’s intended use. If we know which task we are solving with a model, however, we can instead assess the abstracted model directly and evaluate the actions it recommends for this task [Kinney and Watson, 2020]. We need not require that the model approximates some reference model as long as it is useful to solve the given task. If, instead, we do not know what the abstracted model is to be deployed for, we need to revert to an error measure that rates models by their resemblance of a reference model such that we can select models that are useful for a wide range of tasks. Our information theoretic error measure evades the arbitrary choice of a metric on the outcome space.

1.6 A CATEGORICAL APPROACH AND COMPOSITIONAL TRANSFORMATIONS

We propose *compositionality* as a desideratum for notions of error for transformations between causal models (cf. Section 1.4). In our framework, the error of a composite transformation can be bounded in terms of the error of the component transformations: if $f : M \rightarrow M'$ and $g : M' \rightarrow M''$ are transformations between causal models and e, e' are the corresponding errors, then the composite transformation $g \circ f : M \rightarrow M''$ exists and has error at most $e + e'$.

We propose the language of *category theory* to discuss abstractions with a focus on compositionality. Category theory studies mathematical objects in terms of the compositional transformations between them. A *category* \mathcal{C} consists of

1. a collection of *objects*, usually written $\text{ob } \mathcal{C}$,
2. for each pair of objects X, Y , a collection of *morphisms* $f : X \rightarrow Y$,
3. a notion of associative and unital *composition*, assigning to each pair of morphisms $f : X \rightarrow Y, g : Y \rightarrow Z$ a composite $g \circ f : X \rightarrow Z$, and
4. for each object X , an identity morphism $1_X : X \rightarrow X$.

For example, there is a category where the objects are vector spaces, the morphisms are linear functions, and composition is composition of functions.

A category is an abstract mathematical structure, where “object” and “morphism” label two parts of that structure. For example, we can define a category with two objects $\{\star, \clubsuit\}$, the identity morphisms $1_\star : \star \rightarrow \star$ and $1_\clubsuit : \clubsuit \rightarrow \clubsuit$, and one morphism $f : \star \rightarrow \clubsuit$. The notation $f : X \rightarrow Y$ is overloaded and implies that f is some morphism between the objects X and Y , rather than a *function* f from X to Y .

To establish an interdisciplinary meeting ground, we introduce the categorical concepts on a level of detail necessary to understand and gain intuition about our framework. See, for example, MacLane [1971] or Riehl [2016], for comprehensive introductions to category theory.

2 CAUSAL MODELS

To ease the exposition of our categorical framework, we briefly recap the conventional introduction of causal models and their key properties in Section 2.1. In Section 2.2, we introduce a category of finite interventional causal models, where the objects are causal models over finitely many variables with finite outcome spaces. The morphisms in this category are model transformations with an associated error which we formalise via an enriched category and prove to be compositional in Section 2.3.

2.1 RECAP: STRUCTURAL CAUSAL MODELS

For context, we briefly introduce Structural Causal Models (SCMs). For details, we refer to, among others, Spirtes et al. [2001], Pearl [2009], Bollen [1989], Peters et al. [2017].


Definition 2.1 (Structural Causal Model (SCM)). *Let \mathbb{I} be an index set. Let $E = (E_i : i \in \mathbb{I})$ be a collection of independent variables with distribution $\mathbb{P}_E = \otimes_{i \in \mathbb{I}} \mathbb{P}_{E_i}$. Let \mathbf{S} be a set of structural equations*

$$X_i = f_i(X_1, \dots, X_{i-1}, E_i)$$

for $f_i : \prod_{j=1}^{i-1} \mathbf{X}_j \times \mathbf{E}_i \rightarrow \mathbf{X}_i$ measurable functions and $i \in \mathbb{I}$. Let \mathbf{I} be a set of interventions which we denote by $\text{do}(X_{k_1} = x_{k_1}, \dots, X_{k_l} = x_{k_l})$ for $k_1, \dots, k_l \subseteq [d]$ and $x_{k_j} \in \mathbf{X}_{k_j}$ and which identify some structural equations in \mathbf{S} to be replaced by the equations $\mathbf{X}_{k_j} = x_{k_j}$, $j \in [l]$.

We call $(\mathbf{S}, \mathbf{I}, \mathbb{P}_E)$ a structural causal model.

SCMs induce sets of distributions over $X = (X_i : i \in \mathbb{I})$. The distribution \mathbb{P}_X^\emptyset induced by \mathbb{P}_E and the structural equations \mathbf{S} is called the observational distribution. Further distributions $\mathbb{P}_X^{\text{do}(i)}$ are obtained for each intervention $i \in \mathbf{I}$ by changing the respective structural equations according to i and considering the distribution induced by \mathbb{P}_E and this new set of structural equations. Thus, an SCM can be understood as a structured set of joint distributions over X where the distributions are indexed by interventions. The mental picture may be depicted as



The distributions model the system under different interventions that force certain variables to take on certain values. The partial ordering of the distributions reflects the compositionality of interventions, which is crucial for causally consistent reasoning between two models [Rubenstein et al.,

2017]. Acyclicity is a common assumption and ensures that observational and interventional distributions are well-defined; cyclic models are intricate [Bongers et al., 2018].

2.2 A CATEGORY OF FINITE INTERVENTIONAL CAUSAL MODELS

We begin by defining the objects of our category FinMod of finite interventional causal models:

Definition 2.2 (Finite interventional causal model). A finite interventional model $M = (G^M, \mathbf{X}^M, \varphi)$ consists of

1. a finite directed acyclic graph G^M , called causal graph of M , with vertices called variables $V(M)$ of M ,
2. for each variable v , a finite set \mathbf{X}_v^M of possible values of v , and
3. for each variable v , a Markov kernel called mechanism

$$\varphi_v^M : \prod_{v' : v' \rightarrow v \text{ in } G^M} \mathbf{X}_{v'}^M \rightarrow \mathbf{X}_v^M .^1$$

For each root node v , there is a kernel $\varphi_v^M : * \rightarrow \mathbf{X}_v^M$, that is, a probability distribution on \mathbf{X}_v^M .

A finite interventional causal model induces distributions:

Definition 2.3 (Interventional distributions). Given a subset $S \subseteq V(M)$ of the variables in a model M and corresponding values $x_v, v \in S$, there is a well-defined interventional distribution, a kernel

$$I_S : \prod_{v \in S} \mathbf{X}_v^M \rightarrow \prod_{v \in V(M)} \mathbf{X}_v^M,$$

determined by the condition that for $v \in S$, $I_S(x)_v = x_v$ with probability 1, and the conditional distribution of each other variable $v \notin S$, given its parents, is given by the mechanism φ_v^M .

When S is empty, $\prod_{v \in S} \mathbf{X}_v^M = *$, we obtain the observational distribution as the joint distribution over all variables under the null intervention.


We may write $P(- \mid \text{do}(v = x_v, v \in S))$ for $I_S((x_v)_{v \in S})$.

Acyclicity of the causal graph ensures that the distributions in Definition 2.3 are well-defined. To establish the parallels between the categorical and the classical perspective, we (re-)prove this result using string diagram surgery in Appendix B. String diagrams are widely used in category theory to depict constructions in *monoidal categories* such

¹In other words, a function $\prod_{v' \rightarrow v} \mathbf{X}_{v'}^M \rightarrow \Delta(\mathbf{X}_v^M)$.

Since we do not consider counterfactuals in the present article, it is sufficient to specify these kernels instead of functional equations and distributions on the exogenous variables.

as FinStoch [Fritz, 2020]. The proofs rely on rewiring diagrams such as the following (read bottom-to-top)



without changing the resulting distribution; above diagram, for example, depicts a kernel $\mathbf{X}_X^M \rightarrow \mathbf{X}_y^M \times \mathbf{X}_X^M \times \mathbf{X}_{y'}^M$, informally described as “given $x \in \mathbf{X}_X^M$, sample $y \in \mathbf{X}_y$ from the distribution $\varphi_y^M(x)$ and independently sample y' from the distribution $\varphi_{y'}^M(x)$, then return the tuple (y, x, y') ”.

Next, we define the morphisms in FinMod :

Definition 2.4 (Model transformation). A transformation of models $f : M \rightarrow M'$ consists of

1. a surjective vertex map $f_V : V(M) \rightarrow V(M')$,
2. for each $v \in G^{M'}$, a measurement function

$$f_v^m : \prod_{f_V(v')=v} \mathbf{X}_{v'}^{M'} \rightarrow \mathbf{X}_v^{M'}, \text{ and}$$

3. for each $v \in G^{M'}$, an intervention kernel

$$f_v^i : \mathbf{X}_v^{M'} \rightarrow \prod_{f_V(v')=v} \mathbf{X}_{v'}^M.$$

A model transformation is interpreted as follows:

1. Each high-level¹ variable $v \in V(M')$ in M' abstracts over a set $f_V^{-1}(\{v\})$ of low-level variables in M .
2. The high-level observation of v is determined by the values $x_{v'}$ of the low-level variables v' via $f_v^m(x_{v'})$.
3. For each intervention $\text{do}(v = x_v)$ on high-level variables v , there is a *distribution* $f_v^i(x_v)$ of corresponding interventions on the low-level variables $f_V^{-1}(\{v\})$.

To sum up, f_V^m is a map from low- to high-level measurements, while f_V^i is a map from high- to low-level interventions.

This notion of model transformation satisfies the desiderata:

- Variables in the high-level model are explicitly defined relative to the reference model (cf. Section 1.1).

¹When $f : M \rightarrow M'$, we may, for enhanced intuition, think of the models M and M' as “low-level” and “high-level”, respectively. It is not required that M' be more high-level than M .

- The relation of high-level interventions to interventions in the reference model is explicit (cf. Section 1.2).
- The intervention kernel explicitly captures the uncertainty in the low-level implementation of high-level interventions, enhancing transparency and allowing for the error due to intervention ambiguity to be part of our error when assessing the abstraction: if a high-level intervention is related to low-level interventions with substantially different effects we expect a large approximation error (cf. Section 1.3).

The last point bears additional explanation. Rubenstein et al. [2017] argue for restricting the intervention set on the low-level to enable any simplification at all when moving to a higher-level model. While this approach is also transparent about which content of a causal model is to be abstracted away, the generalisation proposed by Beckers et al. [2019] offers a practical advantage. While “setting the temperature to t ” in a high-level model may be implemented by multiple configurations of all gaseous particles’ velocities, not all low-level implementations are possible or equiprobable. It may be impossible, for instance, that the temperature be raised by imparting an absurdly high velocity to a single molecule while leaving the others unchanged. Following Rubenstein et al. [2017], we may remove such interventions from the set of valid interventions. Yet, their approach cannot encode that among all possible low-level configurations with the same average velocity, some are more probable than others; especially if only certain actionable interventions, such as setting several ACs in a room, are considered on the higher-level. We therefore follow Beckers et al. [2019] and demand high-level interventions to be linked to low-level interventions by an intervention kernel instead of a one-to-many mapping.

2.3 COMPOSITIONAL ERROR

We have not yet imposed any compatibility between the distributions induced by the high-level and the low-level causal model. We develop the notion of transformation error to reflect the level of agreement between two models with a morphism between them. The notion depends on the Jensen-Shannon divergence and the category FinStoch of kernels between finite sets. All proofs can be found in Appendix C.

Definition 2.5 (FinStoch (see also Fritz [2020])). *Let FinStoch be the category where*

1. *objects are finite sets,*
2. *a morphism f from \mathbf{X} to \mathbf{Y} is a kernel, that is, a map from \mathbf{X} to the set $\Delta(\mathbf{Y})$ of probability distributions on \mathbf{Y} ,² and*

²One can think of a morphism $f : \mathbf{X} \rightarrow \mathbf{Y}$ in FinStoch as a stochastic matrix where entries reflect the probabilities to transition from $x \in \mathbf{X}$ to $y \in \mathbf{Y}$.

3. *composition is by integration: if $f : \mathbf{X} \rightarrow \mathbf{Y}$ and $g : \mathbf{Y} \rightarrow \mathbf{Z}$ are kernels, their composition $g \circ f$ is*

$$(g \circ f)(x) = \int_{y \in \mathbf{Y}} g(y) d(f(x))(y) \in \Delta(\mathbf{Z}).$$

$\text{FinStoch}(\mathbf{X}, \mathbf{Y})$ *denotes the set of kernels from \mathbf{X} to \mathbf{Y} .*

To measure the error introduced by a causal model transformation, we define a distance between probability measures based on the *Jensen-Shannon divergence* (JSD) [Endres and Schindelin, 2003]:

Definition 2.6 (Jensen-Shannon divergence (JSD)). *Let p_0, p_1 be distributions on a finite set \mathbf{X} . Let B be a random variable with $P(B = 0) = P(B = 1) = 1/2$, that is, a fair coin flip. Let X be a random variable, valued in \mathbf{X} , with $P(X = x \mid B = i) = p_i(x)$. The Jensen-Shannon divergence $\text{JSD}(p_0, p_1)$ is defined as the mutual information between B and X .*

Intuitively, JSD answers the following question: If we learn the value x of X , sampled either from p_0 or p_1 , how much information does $X = x$ reveal about which of the two distributions it was sampled from? Based on the JSD, we define a distance between probability measures as follows:

Definition 2.7 (Jensen-Shannon distance). *For $f, g \in \text{FinStoch}(\mathbf{X}, \mathbf{Y})$, the Jensen-Shannon distance is defined as*

$$d_{\text{JSD}}(f, g) = \sup_{x \in \mathbf{X}} \sqrt{\text{JSD}(f(x), g(x))} \leq 1.$$

d_{JSD} *is a metric on $\text{FinStoch}(\mathbf{X}, \mathbf{Y})$.*

Proposition 2.8 (kernel composition is short). *The composition of kernels*

$$\text{FinStoch}(\mathbf{X}, \mathbf{Y}) \otimes \text{FinStoch}(\mathbf{Y}, \mathbf{Z}) \rightarrow \text{FinStoch}(\mathbf{X}, \mathbf{Z})$$

is a short map, that is, for any $f_1, f_2 \in \text{FinStoch}(\mathbf{Y}, \mathbf{Z})$, $g_1, g_2 \in \text{FinStoch}(\mathbf{X}, \mathbf{Y})$ it holds that

$$d_{\text{JSD}}(f_1 \circ g_1, f_2 \circ g_2) \leq d_{\text{JSD}}(f_1, f_2) + d_{\text{JSD}}(g_1, g_2).$$

Remark 2.9. The above can be summarized as follows: d_{JSD} defines an *enrichment* of FinStoch in the monoidal category Met of metric spaces. We provide a brief description of enriched categories in Appendix D.

In the following lemma we prove that JSD is compositional, which is key for the compositionality of our notion of abstraction error. JSD is only an exemplary choice of distance. More precisely, causal model transformation error can be defined analogously and its compositionality is guaranteed by Proposition 2.12 also for any other distance for which we can prove analogs of Proposition 2.8 and Lemma 2.10.

We use diagrams to depict some collection of objects and morphisms in a category. We encourage readers unfamiliar with these diagrams to understand “consider a diagram ...” as “consider a collection of finite sets and kernels ...”.

Lemma 2.10 (JSD is compositional). *Consider a diagram (not necessarily commutative³) in FinStoch of the following form:*

$$\begin{array}{ccc}
 A & \xrightarrow{f} & A' \\
 \downarrow a & & \uparrow b' \\
 B & \xrightarrow{g} & B' \\
 \downarrow b & & \uparrow c' \\
 C & \xrightarrow{h} & C'
 \end{array} \quad (2.1)$$

Then $d_{\text{JSD}}(f, b'c'hab) \leq d_{\text{JSD}}(f, b'ga) + d_{\text{JSD}}(g, c'hb)$.

Visually, we imagine first replacing the morphism f with $b'ga$, incurring error $d_{\text{JSD}}(f, b'ga)$, then replacing the morphism g with $c'hb$, incurring error $d_{\text{JSD}}(g, c'hb)$. Proposition 2.8 ensures that the alteration of one part of a composition does not create more error than the error associated with the alteration itself. The triangle inequality ensures that successive alterations combine in a natural way.

We define our notion of compositional transformation error:

Definition 2.11 (Transformation error). *Let $f : M \rightarrow M'$ be a transformation of models in FinMod, and $S \subseteq V(M')$ a subset of variables.*

The error associated with S is the Jensen-Shannon distance $d_{\text{JSD}}(I_{\{v_i\}}^{M'}, f^m \circ I_{f_V^{-1}(\{v_i\})}^M \circ f^i)$, which reflects the failure of the following diagram to commute

$$\begin{array}{ccc}
 \prod_{i \in S} \mathbf{X}_{v_i}^{M'} & \xrightarrow{\quad} & \prod_{v \in V(M')} \mathbf{X}_v^{M'} \text{ high} \\
 \downarrow \text{intervention} & & \uparrow \text{distribution} \\
 \prod_{i \in S} \prod_{f(v)=v_i} \mathbf{X}_v^M & \xrightarrow{\quad} & \prod_{v \in V(M)} \mathbf{X}_v^M \text{ low}
 \end{array}$$

The error of f is the maximal error associated with any subset of $V(M')$:

$$\text{error}(f) = \max_{S \subseteq V(M')} d_{\text{JSD}}(I_{\{v_i\}}^{M'}, f^m \circ I_{f_V^{-1}(\{v_i\})}^M \circ f^i)$$

The maximum exists, since $V(M')$ is finite.

The interpretation is as follows: We capture how different the distribution in the high-level model is compared to picking a corresponding low-level intervention (f^i), considering its implementation in M ($I_{f_V^{-1}(\{v_i\})}^M$), and measuring on the high-level (f^m).

We prove that this notion of error is compositional:

³A diagram of this type is *commutative* if each way of composing the depicted morphisms yields the same result – in this case, if $f = b'ga$ and $g = c'hb$.

Proposition 2.12 (Transformation error is compositional). *Let $f : M \rightarrow M', g : M' \rightarrow M''$ be transformations between models in FinMod.*

Then $\text{error}(fg) \leq \text{error}(f) + \text{error}(g)$.

Proof. Let $\{v_i = x_i\}$ be any intervention in M'' , and consider this diagram:

$$\begin{array}{ccc}
 \prod_{i \in I} \mathbf{X}_{v_i}^{M''} & \xrightarrow{\quad} & \prod_{v \in V(M'')} \mathbf{X}_v^{M''} \text{ high} \\
 \downarrow & & \uparrow \\
 \prod_{i \in I} \prod_{f(v)=v_i} \mathbf{X}_v^{M'} & \xrightarrow{\quad} & \prod_{v \in V(M')} \mathbf{X}_v^{M'} \text{ mid} \\
 \downarrow & & \uparrow \\
 \prod_{i \in I} \prod_{f(v)=v_i} \prod_{g(v')=v} \mathbf{X}_{v'}^M & \xrightarrow{\quad} & \prod_{v \in V(M)} \mathbf{X}_v^M \text{ low} \\
 \text{intervention} & & \text{distribution}
 \end{array}$$

By assumption, the failure of the top diagram to commute is $\leq \text{error}(f)$ and that of the bottom diagram is $\leq \text{error}(g)$, so the failure of the composite to commute is $\leq \text{error}(f) + \text{error}(g)$. Since this holds for an arbitrary intervention, we have $\text{error}(fg) \leq \text{error}(f) + \text{error}(g)$. \square

2.4 NOTIONS OF ABSTRACTION AND ERROR

For the Jensen-Shannon distance

$$d_{\text{JSD}} : \text{FinStoch}(\mathbf{X}, \mathbf{Y})^2 \rightarrow [0, \infty]$$

the following are the key properties for our development of a compositional account of causal model transformations:

reflexivity $d_{\text{JSD}}(f, f) = 0$

triangle inequality for any kernels $f_1, f_2, f_3 : \mathbf{X} \rightarrow \mathbf{Y}$, we have $d_{\text{JSD}}(f_1, f_3) \leq d_{\text{JSD}}(f_1, f_2) + d_{\text{JSD}}(f_2, f_3)$

compositionality for any $g_1, g_2 : \mathbf{X} \rightarrow \mathbf{Y}, f : \mathbf{Y} \rightarrow \mathbf{Z}$ and $h : \mathbf{W} \rightarrow \mathbf{X}$, we have $d_{\text{JSD}}(fg_1, fg_2) \leq d_{\text{JSD}}(g_1, g_2)$ and $d_{\text{JSD}}(g_1h, g_2h) \leq d_{\text{JSD}}(g_1, g_2)$

It is reasonable that $d_{\text{JSD}}(f, g) = 0 \Rightarrow f = g$, which ensures that if there is *no* error, the two distributions in question can be considered indistinguishable. This also rules out pathological distances with $d(f, g) = 0$ for all f, g .

It is fruitful to discuss possible relaxations of the underlying distance when developing a compositional framework such as ours. For example, the symmetry $d_{\text{JSD}}(f, g) = d_{\text{JSD}}(g, f)$ is not essential. The interpretation of any chosen distance $d(f, g)$ in the definition of an error measure is ‘‘How bad is it to predict f when the true distribution is g ?’’. For this, it may be reasonable for the underlying notion to be asymmetric. It may also be reasonable to replace the triangle inequality with modified versions like

$d(f, h)^p \leq d(f, g)^p + d(g, h)^p$ for some $p > 0$. For example, one could use the Jensen-Shannon divergence directly, which satisfies this inequality for $p = 1/2$. The key requirement for compositionality is a useful bound on $d(f, h)$ in terms of $d(f, g)$ and $d(g, h)$.

The compositionality requirement on the distance may not be obvious. Indeed, we may expect most natural notions of distance to satisfy it. For a counterexample, consider variables valued in metric spaces and a Wasserstein distance. Here, if one allows any kernel between metric spaces, compositionality fails: if $\mathbf{X} = \{1, 2\}$, $d_{\mathbf{X}}(1, 2) = 1$, $f_i : * \rightarrow \mathbf{X}$ is the point-distribution at i , then the sup-Wasserstein distance of f_1 and f_2 is 1. If we post-compose this with the identity map $i : \mathbf{X} \rightarrow \mathbf{Y}$, where $Y = \{1, 2\}$, $d_{\mathbf{Y}}(1, 2) = 2$, then $d_{\mathbf{Y}}(if_1, if_2) = 2 \not\leq 1 = d_{\mathbf{X}}(f_1, f_2)$. This is commonly avoided by restricting to *short* maps, that is, to those kernels $g : \mathbf{X} \rightarrow \mathbf{Y}$ such that, for each $x, x' \in \mathbf{X}$, $d_{\mathbf{Y}}(g(x), g(x')) \leq d_{\mathbf{X}}(x, x')$.

It is possible to replace JSD with another metric. For example, the metric considered by [Beckers et al., 2019] is induced by a chosen metric on the values of the random variables. In our terms, this amounts to replacing d_{JSD} with the 1-Wasserstein (or Kantorovich) distance, and replacing finite sets with a category of metric spaces. One can show that the desired compositionality property holds as long as the maps between spaces are required to be short (cf. Proposition 2.8 and see, for example, Fritz and Perrone [2017]). This shortness requirement is not surprising: if the map $f : X \rightarrow Y$ maps two points that are indistinguishable to two completely different points, we cannot expect that this map preserve distances between distributions. One could develop an analogous version of our theory for metric spaces; in favour of a clear and concise exposition of our conceptual contribution we refrain from this development here.

3 FUNCTORS ON ABSTRACTIONS

Until now, we have considered the problem of comparing two models and developed a compositional approach to measuring the error incurred by replacing one model by another. The discussion so far has considered this error *relative to the original model*: for a transformation $f : M \rightarrow M'$; we compare the prediction made by M' to the prediction made by M and translated by f into a prediction about the variables in M' . This analysis elides the fact that both M and M' are imperfect approximations of reality. In fact, M' may be closer to reality than M even if $\text{error}(f) > 0$.

In most cases, however, modelling does not start from a ground-truth model, which we then seek to approximate. Instead, we may have different models for the same real-world system that are related to each other by transformations. The errors of each model may be empirically estimated by performing experiments. Here we discuss what can be said

about the “true” error of some M' in terms of the “true” error of another M and $\text{error}(f)$ of $f : M \rightarrow M'$.

We make the following informal definition: An *implemented model* M consists of a finite interventional causal model (also called M), some specified procedure for obtaining a measurement of the variables in the model, and a way of physically implementing each intervention in the model.

The interpretation of the above definitions is that an implementation of a model is an imaginary transformation from a somewhat “idealized model” representing “ground truth”. We can capture how well this implemented model describes reality by an operational definition of error in terms of a two-player recognition game:

1. First, player A chooses some intervention i , which is shown to player B.
2. Then, according to a fair coin flip, player A either physically implements the given intervention and measures the variables in real life, or they sample the variable values according to the interventional distribution $P_M(- \mid \text{do}(i))$ described by the finite interventional causal model M .
3. The outcome of step 2. (but not of the coin flip) is also revealed to player B. Player B must now choose a subjective probability $p \in [0, 1]$ that the measurement was taken in real life.
4. If the measurement was real, player B scores $1 - \log_2 p$, else they score $1 - \log_2(1 - p)$.

The optimal expected score for player B is obtained by choosing for p the conditional probability of the measurement being taken in real given the variable values; the optimal expected score is the mutual information between the coin flip and the values revealed to player B, that is, the Jensen-Shannon distance between real and model distribution Gneiting and Raftery [2007] In principle, this recognition game allows to empirically assess the error of an implemented model through interventional experimentation.

Let M, M' be two implemented models and define an *implementation-preserving transformation* as a model transformation $f : M \rightarrow M'$ satisfying both of the following conditions:

1. For each intervention i in M' , its implementation has the same effect as choosing an intervention in M according to the distribution $f^i(i)$ and implementing that.
2. Measuring a variable x in M' is the same as measuring $f_V^{-1}(\{x\})$ in M and then applying f^m .

Here, the terms “has the same effect” and “the same” are to be informally understood. We assume, for example, that the experimenter can sample a random intervention from $f^i(i)$ in a way that does not interfere with the experiment itself.

Based on this, we might expect the following to hold:

- Given an implemented model M and a transformation $f : M \rightarrow M'$ to some model M' , there is a unique implementation of M' such that f is implementation-preserving.
- Given an implementation-preserving transformation $f : M \rightarrow M'$, the error of M' is at most $\text{error}(f)$ greater than the error of M .

Category theorists may recognise this as a *functor*: The assignment to each abstract model M of its set of implementations and their corresponding error is a functor from FinMod to a category of error spaces (cf. Appendix D).

The “ground truth” functor above is inaccessible to mathematical description, but we can consider variants of this idea. For example, if M is a model, the assignment

$$M' \mapsto (\{f : M \rightarrow M'\}, f \mapsto \text{error}(f) + \epsilon)$$

is a functor. This corresponds to a situation where we know that reality is described by the model M with error at most ϵ , while we are without access to ground truth itself. Nonetheless, we can evaluate models by only comparing them to M ; we have to add an extra ϵ of error to account for the fact that M itself is an imperfect approximation of reality.

The operational meaning is: If a transformation $M \rightarrow M'$ has error ϵ , then the expected score of player B in the recognition game for M' can be bounded in terms of the expected score of player B in the recognition game for M and ϵ , namely $\sqrt{e(M')} \leq \sqrt{e(M)} + \epsilon$. This way, the error of a transformed model in describing reality can be bounded without additional interventional real-world experiments.

The above analysis hinges on the compositionality of abstraction error, such that the error $\text{error}(f)$ provides a bound on how much the error of the implemented model M' is increased compared to M . This underlines the fruitfulness of compositionality as a desideratum and of the categorical abstract viewpoint: We can reasonably talk about bounding the error of a causal model relative to ground truth as long as we have a reference model M for which we evaluated the error by the intervention procedure outlined above.

4 DISCUSSION

We provided a categorical perspective on causal model transformations. Our approach is based on a category of finite interventional causal models and satisfies an important desideratum: compositionality of model transformations and the associated approximation errors. While we consider a category of causal models and transformations between them, existing work on the application of category theory to the domain of causal modelling has studied one causal model via categorical tools. Fong [2013], for example, develops the theory of directed acyclic graph models using syntactical categories. This is further developed by Jacobs

et al. [2019], who demonstrate how to carry out causal inference by string diagram manipulations. Fritz [2020] lays out the foundations for an ambitious programme of developing probability theory in the language of categories. Working in this framework, Patterson [2020] develops an analogy between statistics and universal algebra, where a statistical model becomes a model of a theory, in the sense of logic. This separation of theory and model is akin to the approach presented by Bongers et al. [2018]: detaching the structural equations (the *theory*) from the random variables that simultaneously (almost surely) solve those equations (the *model*), they provide a measure theoretic treatment of cyclic models.

Conceptualising a framework for causal model transformations can be motivated from different vantage points. First, it helps characterise when observable variables may be ill-suited for a causal description by viewing the observables as a transformation of underlying causal entities with high error. For example, in the analysis of electroencephalographic data we may wish to recover signals that correspond to cortical activity instead of reasoning about interventions on mixed electrode signals [Weichwald et al., 2016]. Second, if observables are ill-suited for a causal description, we may wish to find a transformation that yields variables amenable to a causal description. Chalupka et al. [2015], for example, present how to learn the macroscopic visual cause of some behaviour from observed pixel values. Third, we may be interested in abstracting or aggregating information to obtain a macro-level description of a system that is pragmatically more useful as it represents the information necessary for a certain task more clearly than a complex fine-grained model [Hoel et al., 2013, Hoel, 2017, Kinney and Watson, 2020, Weichwald, 2019]. Last, approaches to infer causality between latent causal variables based on observed variables or time-subsampled observations may be embedded within a framework of causal model transformation where transformations encode which variables or time-points are unobserved [Hytinen et al., 2016, Silva et al., 2006].

Our category theoretic framework of causal model transformations is instructive to clarify the assumptions and arguments required to proof its compositionality: We require the distance between kernels to be compatible with composition of kernels, that is, beyond the triangle inequality we require analogs of Proposition 2.8 and Lemma 2.10. This condition is natural from a category-theoretical perspective. The formal tools of category theory enable diagrammatic reasoning and a simple proof that the resulting framework of causal model transformations and their abstraction errors is compositional.

Acknowledgements

We thank the anonymous reviewers for their constructive comments that helped improve the interdisciplinary presentation. SW was supported by the Carlsberg Foundation.

References

- Sander Beckers and Joseph Y Halpern. Abstracting Causal Models. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*. AAAI Press, 2019.
- Sander Beckers, Frederick Eberhardt, and Joseph Y Halpern. Approximate Causal Abstraction. In *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI)*. AUAI Press, 2019.
- Kenneth A Bollen. *Structural Equations with Latent Variables*. John Wiley & Sons, 1989.
- Stephan Bongers, Patrick Forré, Jonas Peters, Bernhard Schölkopf, and Joris M Mooij. Foundations of Structural Causal Models with Cycles and Latent Variables. *arXiv preprint arXiv:1611.06221*, 2018.
- Krzysztof Chalupka, Pietro Perona, and Frederick Eberhardt. Visual Causal Feature Learning. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence (UAI)*. AUAI Press, 2015.
- Frederick Eberhardt. Green and grue causal variables. *Synthese*, 193(4):1029–1046, 2016.
- Dominik M. Endres and Johannes E. Schindelin. A new metric for probability distributions. *IEEE Transactions on Information Theory*, 49(7):1858–1860, 2003.
- Brendan Fong. Causal Theories: A Categorical Perspective on Bayesian Networks. *arXiv preprint arXiv:1301.6201*, 2013.
- Tobias Fritz. A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics. *Advances in Mathematics*, 370:107239, 2020.
- Tobias Fritz and Paolo Perrone. A Probability Monad as the Colimit of Spaces of Finite Samples. *arXiv preprint arXiv:1712.05363*, 2017.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Erik P Hoel. When the Map Is Better Than the Territory. *Entropy*, 19(5), 2017.
- Erik P Hoel, Larissa Albantakis, and Giulio Tononi. Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences*, 110(49):19790–19795, 2013.
- Antti Hyttinen, Sergey Plis, Matti Järvisalo, Frederick Eberhardt, and David Danks. Causal discovery from subsampled time series data by constraint optimization. In *Proceedings of the Eighth International Conference on Probabilistic Graphical Models*, volume 52 of *Proceedings of Machine Learning Research*, pages 216–227. PMLR, 2016.
- Bart Jacobs, Aleks Kissinger, and Fabio Zanasi. Causal inference by string diagram surgery. *Foundations of Software Science and Computation Structures*, pages 313–329, 2019.
- Dominik Janzing, Paul Rubenstein, and Bernhard Schölkopf. Structural causal models for macro-variables in time-series. *arXiv preprint arXiv:1804.03911*, 2018.
- David Kinney and David Watson. Causal feature learning for utility-maximizing agents. In Manfred Jaeger and Thomas Dyhre Nielsen, editors, *Proceedings of the Tenth International Conference on Probabilistic Graphical Models*, volume 138 of *Proceedings of Machine Learning Research*, pages 257–268. PMLR, 2020.
- Saunders MacLane. *Categories for the Working Mathematician*. Springer, 1971. Graduate Texts in Mathematics, Vol. 5.
- Evan Patterson. *The algebra and machine representation of statistical models*. PhD thesis, Stanford University, Department of Statistics, 2020.
- Judea Pearl. *Causality*. Cambridge University Press, 2 edition, 2009.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference*. MIT Press, 2017.
- Emily Riehl. *Category Theory in Context*. Aurora: Dover Modern Math Originals. Dover Publications, 2016. ISBN 9780486820804.
- Paul K Rubenstein, Sebastian Weichwald, Stephan Bongers, Joris M Mooij, Dominik Janzing, Moritz Grosse-Wentrup, and Bernhard Schölkopf. Causal Consistency of Structural Equation Models. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence (UAI)*. AUAI Press, 2017.
- Ricardo Silva, Richard Scheine, Clark Glymour, and Peter Spirtes. Learning the structure of linear latent variable models. *Journal of Machine Learning Research*, 7:191–246, 2006.
- Peter Spirtes and Richard Scheines. Causal Inference of Ambiguous Manipulations. *Philosophy of Science*, 71(5): 833–845, 2004.
- Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2 edition, 2001.
- Sebastian Weichwald. *Pragmatism and Variable Transformations in Causal Modelling*. PhD thesis, ETH Zurich, Zurich, Switzerland, 2019.

Sebastian Weichwald, Moritz Grosse-Wentrup, and Arthur Gretton. MERLiN: Mixture Effect Recovery in Linear Networks. *IEEE Journal of Selected Topics in Signal Processing*, 10(7):1254–1266, 2016.

Preliminary version