

The Causal Frame Problem: An Algorithmic Perspective

Ardavan S. Nobandegani^{1,2} Ioannis N. Psaromiligkos¹

{ardavan.salehinobandegani@mail.mcgill.ca, yannis@ece.mcgill.ca}

¹Department of Electrical & Computer Engineering, McGill University

²Department of Psychology, McGill University

Abstract

The Frame Problem (FP) is a puzzle in philosophy of mind and epistemology, articulated by the Stanford Encyclopedia of Philosophy as follows: “How do we account for our apparent ability to make decisions on the basis only of what is relevant to an ongoing situation without having explicitly to consider all that is not relevant?” In this work, we focus on the causal variant of the FP, the Causal Frame Problem (CFP). Assuming that a reasoner’s mental causal model can be (implicitly) represented by a causal Bayes net, we first introduce a notion called Potential Level (PL). PL, in essence, encodes the relative position of a node with respect to its neighbors in a causal Bayes net. Drawing on the psychological literature on causal judgment, we substantiate the claim that PL may bear on how *time* is encoded in the mind. Using PL, we propose an inference framework, called the PL-based Inference Framework (PLIF), which permits a boundedly-rational approach to the CFP, formally articulated at Marr’s algorithmic level of analysis. We show that our proposed framework, PLIF, is consistent with several findings in the causal judgment literature, and that PL and PLIF make a number of predictions, some of which are already supported by existing findings.

Keywords: Frame Problem; Time and Causality; Bounded Rationality; Algorithmic Level of Analysis; Rational Process Models

1 Introduction

At the core of any decision-making or reasoning task, resides an innocent-looking yet challenging question: Given an inconceivably large body of knowledge available to the reasoner, what constitutes the relevant for the task and what the irrelevant? The question, as it is posed, echoes the well-known Frame Problem (FP) in epistemology and philosophy of mind, articulated by Glymour (1987) as follows: “Given an enormous amount of stuff, and some task to be done using some of the stuff, what is the relevant stuff for the task?” Fodor (1987) comments: “The frame problem goes very deep; it goes as deep as the analysis of rationality.”

The question posed above perfectly captures what is really at the core of the FP, yet, it may suggest an unsatisfying approach to the FP at the algorithmic level of analysis (Marr, 1982). Indeed, the question may suggest the following two-step methodology: In the first step, out of all the body of knowledge available to the reasoner (termed, the model), she has to identify what is relevant to the task (termed, the relevant submodel); it is *only then* that she advances to the second step by performing *reasoning* or *inference* on the identified submodel. There is something fundamentally wrong with this methodology (which we term, *sequential* approach to reasoning) which bears on the following understanding: The relevant submodel, i.e., the portion of the reasoner’s knowledge deemed relevant to the task, oftentimes is so enormous (or even infinitely large) that the reasoner—inevitably bounded

in time and computational resources—would never get to the second step, had she adhered to such a methodology. In other words, in line with the notion of bounded rationality (Simon, 1957), a boundedly-rational reasoner must have the option, if need be, to merely consult a fraction of the potentially large—if not infinitely so—relevant submodel.

Icard and Goodman (2015) elegantly promote this insight when they write: “Somehow the mind must focus in on some “submodel” of the “full” model (including all possibly relevant variables) that suffices for the task at hand and is not too costly to use.”¹ They then ask the following question: “what kind of simpler model should a reasoner consult for a given task?” This is an inspiring question hinting to an interesting line of inquiry as to how to formally articulate a boundedly-rational approach to the FP, at Marr’s (1982) algorithmic level of analysis.

In this work, we focus on the causal variant of the FP, the Causal Frame Problem (CFP), stated as follows: Upon being presented with a causal query, how does the reasoner manage to attend to her causal knowledge relevant to the derivation of the query while rightfully dismissing the irrelevant? We adopt Causal Bayesian Networks (CBNs) (Pearl, 1988; Gopnik et al., 2004, *inter alia*) as a normative model to represent how the reasoner’s internal causal model of the world is structured (i.e., reasoner’s mental model). First, we introduce the notion of Potential Level (PL). PL, in essence, encodes the relative position of a node (representing a propositional variable or a concept) with respect to its neighbors in a CBN. Drawing on the psychological literature on causal judgment, we substantiate the claim that PL may bear on how *time* is encoded in the mind. Equipped with PL, we embark on investigating the CFP at Marr’s algorithmic level of analysis. We propose an inference framework, termed PL-based Inference Framework (PLIF), which aims at empowering the boundedly-rational reasoner to consult (or retrieve²) parts of the underlying CBN deemed relevant for the derivation of the posed query (the relevant submodel) in a *local, bottom-up* fashion until the submodel is fully retrieved. PLIF allows the reasoner to carry out inference at *intermediate* stages of the retrieval process over the thus-far retrieved parts, thereby obtaining lower and upper bounds on the posed causal query.

¹In an informative example on Hidden Markov Models (HMMs), Icard and Goodman (2015) present a setting wherein the relevant submodel is infinitely large—an example which highlights what is wrong with the sequential approach stated earlier.

²The terms “consult” and “retrieve” will be used interchangeably. We elaborate on the rationale behind that in Sec. 5, where we connect our work to Long Term Memory and Working Memory.

We show, in the Discussion section, that our proposed framework, PLIF, is consistent with several findings in the causal judgment literature, and that PL and PLIF make a number of predictions, some of which are already supported by the findings in the psychology literature.

In their work, Icard and Goodman (2015) articulate a boundedly-rational approach to the CFP at Marr’s computational level of analysis, which, as they point out, is from a “god’s eye” point of view. In sharp contrast, our proposed framework PLIF is *not* from a “god’s eye” point of view and hence could be regarded, potentially, as a psychologically plausible proposal at Marr’s algorithmic level of analysis as to how the mind both retrieves and, at the same time, carries out inference over the retrieved submodel to derive bounds on a causal query. We term this *concurrent* approach to reasoning, as opposed to the flawed sequential approach stated earlier.³ The retrieval process progresses in a local, bottom-up fashion, hence the submodel is retrieved *incrementally*, in a *nested* manner.⁴ Our analysis (Sec. 4.1) confirms Icard and Goodman’s (2015) insight that even in the extreme case of having an infinitely large relevant submodel, the portion of which the reasoner has to consult so as to obtain a “sufficiently good” answer to a query could indeed be very small.

2 Potential Level and Time

Before proceeding further, let us introduce some preliminary notations. Random Variables (RVs) are denoted by lower-case bold-faced letters, e.g., \mathbf{x} , and their realizations by non-bold lower-case letters, e.g., x . Likewise, sets of RVs are denoted by upper-case bold-faced letters, e.g., \mathbf{X} , and their corresponding realizations by upper-case non-bold letters, e.g., X . $Val(\cdot)$ denotes the set of possible values a random quantity can take on. Random quantities are assumed to be discrete unless stated otherwise. The joint probability distribution over $\mathbf{x}_1, \dots, \mathbf{x}_n$ is denoted by $\mathbb{P}(\mathbf{x}_1, \dots, \mathbf{x}_n)$. We will use the notation $\mathbf{x}_{1:n}$ to denote the sequence of n RVs $\mathbf{x}_1, \dots, \mathbf{x}_n$, hence $\mathbb{P}(\mathbf{x}_1, \dots, \mathbf{x}_n) = \mathbb{P}(\mathbf{x}_{1:n})$. The terms “node” and “variable” will be used interchangeably. To simplify presentation, we adopt the following notation: We denote the probability $\mathbb{P}(\mathbf{x} = x)$ by $\mathbb{P}(x)$ for some RV \mathbf{x} and its realization $x \in Val(\mathbf{x})$. For conditional probabilities, we will use the notation $\mathbb{P}(x|y)$ instead of $\mathbb{P}(\mathbf{x} = x|\mathbf{y} = y)$. Likewise, $\mathbb{P}(X|Y) = \mathbb{P}(\mathbf{X} = X|\mathbf{Y} = Y)$ for $X \in Val(\mathbf{X})$ and $Y \in Val(\mathbf{Y})$. A generic conditional independence relationship is denoted by $(\mathbf{A} \perp\!\!\!\perp \mathbf{B}|\mathbf{C})$ where \mathbf{A}, \mathbf{B} , and \mathbf{C} represent three mutually disjoint sets of variables belonging to a CBN. Furthermore, throughout the paper, we assume that ε is some negligibly small positive real-valued quantity. Whenever we subtract ε from a quantity, we simply imply a quantity less than but arbitrarily close to the original quantity. The rationale behind adopting such a notation will become clearer in Sec. 4.

³We elaborate more on this in the Discussion section.

⁴The term “nested” implies that the thus-far retrieved submodel is subsumed by every later submodel (provided that the reasoner proceeds with the retrieval process).

Before formally introducing the notion of PL, we articulate in simple terms what the idea behind PL is. PL simply induces a *chronological order* on the nodes of a CBN, allowing the reasoner to encode the timing between cause and effect.⁵ As we will see, PL plays an important role in guiding the retrieval process used in our proposed framework. Next, PL is formally defined, followed by two clarifying examples.

Def. 1. (Potential Level (PL)) Let $par(\mathbf{x})$ and $child(\mathbf{x})$ denote, respectively, the sets of parents (i.e., immediate causes) and children (i.e., immediate effects) of \mathbf{x} . Also let $T_0 \in \mathbb{R} \cup \{-\infty\}$. The PL of \mathbf{x} , denoted by $p_l(\mathbf{x})$, is defined as follows: (i) If $par(\mathbf{x}) = \emptyset$, $p_l(\mathbf{x}) = T_0$, and (ii) If $par(\mathbf{x}) \neq \emptyset$, $p_l(\mathbf{x})$ is a real-valued quantity selected from the interval $(\max_{\mathbf{y} \in par(\mathbf{x})} p_l(\mathbf{y}), \min_{\mathbf{z} \in child(\mathbf{x})} p_l(\mathbf{z}))$ such that $p_l(\mathbf{x}) - \max_{\mathbf{y} \in par(\mathbf{x})} p_l(\mathbf{y})$ indicates the amount of time which elapses between intervening simultaneously on all the RVs in $par(\mathbf{x})$ (i.e., $do(par(\mathbf{x}) = par_x)$) and \mathbf{x} taking its value x in accord with the distribution $\mathbb{P}(x|par_x)$. If $child(\mathbf{x}) = \emptyset$, substitute the upper bound of the given interval by $+\infty$. ■

Parameter T_0 symbolizes the origin of time, as perceived by the reasoner. $T_0 = 0$ is a natural choice, unless the reasoner believes that time continues unboundedly into the past, in which case $T_0 = -\infty$. The next two examples further clarify the idea behind PL. In both examples we assume $T_0 = 0$.




Figure 1: Relation between PL and time. Three hollow dots signify that the depicted CBNs extend into the past and future.

For the first example, let us consider the CBN depicted in Fig. 1(a) containing the RVs \mathbf{x}, \mathbf{y} , and \mathbf{z} with $p_l(\mathbf{x}) = 4, p_l(\mathbf{y}) = 4.7$, and $p_l(\mathbf{z}) = 5$. According to Def. 1, the given PLs can be construed in terms of the relative time between the occurrence of cause and effect as articulated next. Upon intervening on \mathbf{x} (i.e., $do(\mathbf{x} = x)$), after the elapse of $p_l(\mathbf{y}) - p_l(\mathbf{x}) = 0.7$ units of time, the RV \mathbf{y} takes its value y in accord with the distribution $\mathbb{P}(y|x)$. Likewise, upon intervening on \mathbf{y} (i.e., $do(\mathbf{y} = y)$), after the elapse of $p_l(\mathbf{z}) - p_l(\mathbf{y}) = 0.3$ units of time, \mathbf{z} takes its value z according to $\mathbb{P}(z|y)$.

For the second example, consider the CBN depicted in Fig. 1(b) containing the RVs $\mathbf{x}, \mathbf{y}, \mathbf{z}$, and \mathbf{t} with $p_l(\mathbf{x}) = 4, p_l(\mathbf{y}) = 4.7, p_l(\mathbf{z}) = 5$, and $p_l(\mathbf{t}) = 5.6$. Upon intervening on \mathbf{x} (i.e., $do(\mathbf{x} = x)$) the following happens: (i) after the elapse of $p_l(\mathbf{y}) - p_l(\mathbf{x}) = 0.7$ units of time, \mathbf{y} takes its value y according to $\mathbb{P}(y|x)$, and (ii) after the elapse of $p_l(\mathbf{z}) - p_l(\mathbf{x}) = 1$ unit of

⁵More precisely, PL induces a topological order on the nodes of a CBN, with temporal interpretations suggested in Def. 1.

time, \mathbf{z} takes its value z according to $\mathbb{P}(z|x)$. Also, upon intervening simultaneously on RVs \mathbf{y}, \mathbf{z} (i.e., $do(\mathbf{y} = y, \mathbf{z} = z)$), after the elapse of $p_l(\mathbf{t}) - \max_{\mathbf{r} \in \text{par}(\mathbf{t})} p_l(\mathbf{r}) = 0.6$ units of time, \mathbf{t} takes its value t according to $\mathbb{P}(t|y, z)$.

In sum, the notion of PL bears on the underlying time-grid upon which a CBN is constructed, and adheres to Hume’s principle of temporal precedence of cause to effect (Hume, 1748/1975). A growing body of work in psychology literature corroborates Hume’s centuries-old insight, suggesting that the timing and temporal order between events strongly influences how humans induce causal structure over them (Bramley, Gerstenberg, & Lagnado, 2014; Lagnado & Sloman, 2006). The introduced notion of PL is based on the following hypothesis: When learning the underlying causal structure of a domain, humans may as well encode the temporal patterns (or some estimates thereof) on which they rely to infer the causal structure. This hypothesis is supported by recent findings suggesting that people have expectations about the delay length between cause and effect (Greville & Buehner, 2010; Buehner & May, 2004; Schlottmann, 1999). It is worth noting that we could have defined PL in terms of relative *expected time* between cause and effect, rather than relative absolute time. Under such an interpretation, the time which elapses between the intervention on a cause and the occurrence of its effect would be modeled by a probability distribution, and PL would be defined in terms of the expected value of that distribution. Our proposed framework, PLIF, is indifferent as to whether PL should be construed in terms of absolute or expected time. Greville and Buehner (2010) show that causal relations with fixed temporal intervals are consistently judged as stronger compared to those with variable temporal intervals. This finding, therefore, seems to suggest that people expect, to a greater extent, fixed temporal intervals between cause and effect, rather than variable ones—an interpretation which, at least to a first approximation, favors construing PL in terms of relative absolute time (see Def. 1).⁶

3 Informative Example

To develop our intuition, and before formally articulating our proposed framework, let us present a simple yet informative example which demonstrates: (i) how the retrieval process can be carried out in a local, bottom-up fashion, allowing for retrieving the relevant submodel incrementally, and (ii) how adopting PL allows the reasoner to obtain bounds on a given causal query at intermediate stages of the retrieval process.

Let us assume that the posed causal query is $\mathbb{P}(x|y)$ where \mathbf{x}, \mathbf{y} are two RVs in the CBN depicted in Fig. 2(a) with PLs $p_l(\mathbf{x}), p_l(\mathbf{y})$, and let $p_l(\mathbf{x}) > p_l(\mathbf{y})$. The relevant information for the derivation of the posed query (i.e., the relevant submodel) is depicted in Fig. 2(e).

⁶There are cases, however, that, despite the precedence of cause to effect, quantifying the amount of time between their occurrences may bear no meaning, e.g., when dealing with hypothetical constructs. In such cases, PL should be simply construed as a topological ordering. From a purely computational perspective, PL is a generalization of *topological sorting* in computer science.

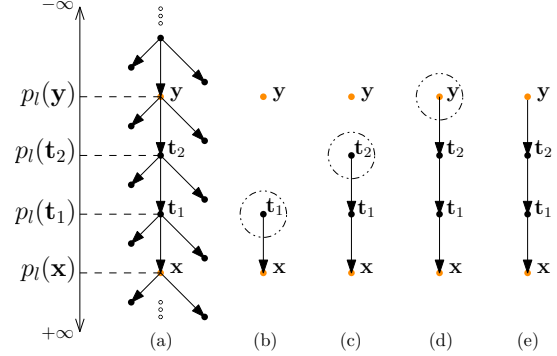


Figure 2: Example. Query variables are shown in orange.

Starting from the target RV \mathbf{x} in the original CBN (Fig. 2(a)) and moving one step backwards,⁷ \mathbf{t}_1 is reached (Fig. 2(b)). Since $p_l(\mathbf{y}) < p_l(\mathbf{t}_1)$, \mathbf{y} must be a non-descendant of \mathbf{t}_1 , and therefore, of \mathbf{x} . Hence, conditioning on \mathbf{t}_1 d -separates \mathbf{x} from \mathbf{y} (Pearl, 1988), yielding $(\mathbf{x} \perp\!\!\!\perp \mathbf{y} | \mathbf{t}_1)$. Thus $\mathbb{P}(x|y) = \sum_{t_1 \in \text{Val}(t_1)} \mathbb{P}(x|y, t_1) \mathbb{P}(t_1|y) = \sum_{t_1 \in \text{Val}(t_1)} \mathbb{P}(x|t_1) \mathbb{P}(t_1|y)$ implying: $\min_{t_1 \in \text{Val}(t_1)} \mathbb{P}(x|t_1) \leq \mathbb{P}(x|y) \leq \max_{t_1 \in \text{Val}(t_1)} \mathbb{P}(x|t_1)$. It is crucial to note that the given bounds can be computed using the information thus-far retrieved, i.e., the information encoded in the submodel shown in Fig. 2(b). Taking a step backwards from \mathbf{t}_1 , \mathbf{t}_2 is reached (Fig. 2(c)). Using a similar line of reasoning to the one presented for \mathbf{t}_1 , having $p_l(\mathbf{y}) < p_l(\mathbf{t}_2)$ ensures $(\mathbf{x} \perp\!\!\!\perp \mathbf{y} | \mathbf{t}_2)$. Therefore, the following bounds on the posed query can be derived, which, crucially, can be computed using the information thus-far retrieved: $\min_{t_2 \in \text{Val}(t_2)} \mathbb{P}(x|t_2) \leq \mathbb{P}(x|y) \leq \max_{t_2 \in \text{Val}(t_2)} \mathbb{P}(x|t_2)$. It is straightforward to show that the bounds derived in terms of \mathbf{t}_2 are equally tight or tighter than the bounds derived in terms of \mathbf{t}_1 . Finally, taking one step backward from \mathbf{t}_2 , \mathbf{y} is reached (Fig. 2(d)) and the exact value for $\mathbb{P}(x|y)$ can be derived, again using the submodel thus-far retrieved (Fig. 2(d)).

We are now ready to present our proposed framework.

4 PL-based Inference Framework (PLIF)

In this section, we intend to elaborate on how, equipped with the notion of PL, a generic causal query of the form⁸ $\mathbb{P}(\mathbf{O} = O | \mathbf{E} = E)$ can be derived where \mathbf{O} and \mathbf{E} denote, respectively, the disjoint sets of target (or *objective*) and observed (or *evidence*) variables. In other words, we intend to formalize how inference over a CBN whose nodes are endowed with PL as an attribute should be carried out. Before we present the main result, a few definitions are in order.

Def. 2. (Critical Potential Level (CPL)) The target variable with the least PL is denoted by \mathbf{o}^* and its PL is referred to as the CPL. More formally, $p_l^* \triangleq \min_{\mathbf{o} \in \mathbf{O}} p_l(\mathbf{o})$ and

⁷Taking one step backwards from variable \mathbf{q} amounts to retrieving all the parents of \mathbf{q} .

⁸We do not consider interventions in this work. However, with some modifications, the presented analysis/results can be extended to handle a generic causal query of the form $\mathbb{P}(\mathbf{O} = O | \mathbf{E} = E, do(\mathbf{Z} = Z))$ where \mathbf{Z} denotes the set of intervened variables.

$\mathbf{o}^* \triangleq \arg \min_{\mathbf{o} \in \mathbf{O}} p_l(\mathbf{o})$. E.g., for the setting given in Fig. 2(a), $\mathbf{o}^* = \mathbf{x}$, and $p_l^* = p_l(\mathbf{x})$. Viewed through the lens of time, \mathbf{o}^* is the furthest target variable into the past, with PL p_l^* .

There are two possibilities: (a) $p_l^* > T_0$, or (b) $p_l^* = T_0$, with T_0 denoting the origin of time; cf. Sec. 2. In the sequel, we assume that (a) holds.⁹

Def. 3. (Inference Threshold (IT) and IT Root Set (IT-RS)) To any real-valued quantity, \mathcal{T} , corresponds a unique set, $\mathbf{R}_{\mathcal{T}}$, obtained as follows: Start at every variable $\mathbf{x} \in \mathbf{O} \cup \mathbf{E}$ with $\text{PL} \geq \mathcal{T}$ and backtrack along all paths terminating at \mathbf{x} . Backtracking along each path stops as soon as a node with PL less than \mathcal{T} is encountered. Such nodes, together, compose the set $\mathbf{R}_{\mathcal{T}}$. It follows that: $\max_{\mathbf{t} \in \mathbf{R}_{\mathcal{T}}} p_l(\mathbf{t}) < \mathcal{T}$. \mathcal{T} and $\mathbf{R}_{\mathcal{T}}$ are termed, respectively, Inference Threshold (IT) and the IT Root Set (IT-RS) for \mathcal{T} .

For example, the set of variables circled at the stages depicted in Figs. 2(b-d) are the IT-RSs for $\mathcal{T} = p_l(\mathbf{x}) - \epsilon$, $\mathcal{T} = p_l(\mathbf{t}_1) - \epsilon$, and $\mathcal{T} = p_l(\mathbf{t}_2) - \epsilon$, respectively. Note that instead of saying $\mathcal{T} = p_l(\mathbf{x}) - \epsilon$ we could have said: for any $\mathcal{T} \in (p_l(\mathbf{t}_1), p_l(\mathbf{x}))$. However, expressing ITs in terms of ϵ liberates us from having to express them in terms of intervals, thereby simplifying the exposition. We would like to emphasize that the adopted notation should not be construed as implying that the assignment of values to ITs is such a sensitive task that everything would have collapsed, had IT not been chosen in such a fine-tuned manner. To recap, in simple terms, \mathcal{T} bears on how far into the past a reasoner is consulting her mental model in the process of answering a query, and $\mathbf{R}_{\mathcal{T}}$ characterizes the furthest-into-the-past concepts entertained by the reasoner in that process.

Next, we formally present the main idea behind PLIF, followed by its interpretation in simple terms.

Lemma 1. Let $\mathbb{P}(O|E)$ denote the posed causal query, with \mathbf{O} and \mathbf{E} denoting, respectively, the disjoint sets of target and observed variables. For any chosen IT $\mathcal{T} < p_l^*$ and its corresponding $\mathbf{R}_{\mathcal{T}}$, define $\mathbf{S} \triangleq \mathbf{R}_{\mathcal{T}} \setminus \mathbf{E}$. Then the following holds: $\min_{S \in \text{Val}(\mathbf{S})} \mathbb{P}(O|S, E) \leq \mathbb{P}(O|E) \leq \max_{S \in \text{Val}(\mathbf{S})} \mathbb{P}(O|S, E)$. Crucially, the provided bounds can be computed using the information encoded in the submodel retrieved in the very process of obtaining the $\mathbf{R}_{\mathcal{T}}$. \square

The message of Lemma 1 is quite simple: For any chosen inference threshold \mathcal{T} which is further into the past than \mathbf{o}^* , Lemma 1 ensures that the reasoner can condition on \mathbf{S} and obtain the reported lower and upper bounds on the query by using *only* the information encoded in the retrieved submodel.

It is natural to ask under what conditions the exact value to the posed query can be derived using the thus-far retrieved submodel, i.e., the submodel obtained during the identification of $\mathbf{R}_{\mathcal{T}}$. The following remark bears on that.¹⁰

Remark 1. If for IT \mathcal{T} , $\mathbf{R}_{\mathcal{T}}$ satisfies either: (i) $\mathbf{R}_{\mathcal{T}} \subseteq \mathbf{E}$, or (ii) for all $\mathbf{r} \in \mathbf{R}_{\mathcal{T}}$, $p_l(\mathbf{r}) = T_0$, and $\min_{\mathbf{e} \in \mathbf{E}} p_l(\mathbf{e}) > \mathcal{T}$, or

⁹For a discussion on the special case (b), the reader is referred to: <https://arxiv.org/pdf/1701.08100>

¹⁰For a formal proof of Lemma 1, and the rationale behind Remark 1, the reader is referred to: <https://arxiv.org/pdf/1701.08100>

(iii) the lower and upper bound given in Lemma 1 are identical, then the exact value of the posed query can be derived using the submodel retrieved in the process of obtaining $\mathbf{R}_{\mathcal{T}}$. Fig. 2(d) shows a setting wherein (i) and (iii) are both met.

4.1 Case Study

Next, we intend to cast the Hidden Markov Model (HMM) studied in (Icard & Goodman, 2015, p. 2) into our framework. The setting is shown in Fig. 3(left). We adhere to the




Figure 3: Left: The infinite-sized HMM discussed in (Icard & Goodman, 2015) with parameterization: $\mathbb{P}(x_{t+1}|x_t) = \mathbb{P}(\bar{x}_{t+1}|\bar{x}_t) = 0.9$, and $\mathbb{P}(y_t|x_t) = \mathbb{P}(\bar{y}_t|\bar{x}_t) = 0.8$. Right: Applying PLIF on the HMM shown in left. Vertical and horizontal axes denote, respectively, the value of the posed query $\mathbb{P}(x_{t+1}|y_{-\infty:t})$ and the adopted IT \mathcal{T} . The vertical bars depict the intervals within which the query lies due to Lemma 1. The dotted curves—which connect the lower and upper bounds of the intervals—show how the intervals shrink as IT \mathcal{T} decreases.

same parameterization and query adopted therein. All RVs in this section are binary, taking on values from the set $\{0, 1\}$; $\mathbf{x} = x$ indicates the event wherein \mathbf{x} takes the value 1, and $\mathbf{x} = \bar{x}$ implies the event wherein \mathbf{x} takes the value 0. We assume $p_l(\mathbf{x}_{t+i}) = i - 2$.¹¹ We should note that the assignment of the PLs for the variables in $\{\mathbf{y}_{t-i}\}_{i=0}^{+\infty}$ does not affect the presented results in any way. The query of interest is $\mathbb{P}(x_{t+1}|y_{-\infty:t})$. Notice that after performing three steps of the sort discussed in the example presented in Sec. 3 (for the IT $\mathcal{T} = -3 - \epsilon$), the lower bound on the posed query exceeds 0.5 (shown by the red dashed line in Fig. 3(right)). This observation has the following intriguing implication. Assume, for the sake of argument, that we were presented with the following Maximum A-Posterior (MAP) inference problem: Upon observing all the variables in $\{\mathbf{y}_{t-i}\}_{i=0}^{+\infty}$ taking on the value 1, what would be the most likely state for the variable \mathbf{x}_{t+1} ? Interestingly, we would be able to answer this MAP inference problem simply after three backward moves (corresponding to the IT $\mathcal{T} = -3 - \epsilon$). In Fig. 3(right), the intervals within

¹¹Note that the trend of the upper- and lower-bound curves as well as the size of the intervals shown in Fig. 3(right) are insensitive with regard to the choice of PLs for variables $\{\mathbf{x}_{t-i}\}_{i=-1}^{+\infty}$.

which the posed query falls (due to Lemma 1) in terms of the adopted IT \mathcal{T} are depicted.

Our analysis confirms Icard and Goodman’s (2015) insight that even in the extreme case of having infinite-sized relevant submodel (Fig. 3(left)), the portion of which the reasoner has to consult so as to obtain a “sufficiently good” answer to the posed query could happen to be very small (Fig. 3(right)).

5 Discussion

To our knowledge, PLIF is the first inference framework proposed that capitalizes on *time* to constrain the scope of causal reasoning over CBNs, where the term scope refers to the portion of a CBN on which inference is carried out. PLIF does not restrict itself to any particular inference scheme. The claim of PLIF is that inference should be confined within and carried out over retrieved submodels of the kind suggested by Lemma 1 so as to obtain the reported bounds therein. In this light, PLIF can accommodate any inference scheme, including Belief Propagation (BP), and sample-based inference methods using Markov Chain Monte Carlo (MCMC), as two prominent classes of inference schemes. MCMC-based methods have been successful in simulating important aspects of a wide range of cognitive phenomena and accounting for many cognitive biases; cf. (Sanborn & Chater, 2016). Also, work in theoretical neuroscience has suggested mechanisms for how BP and MCMC-based methods could be realized in neural circuits; cf. (Gershman & Beck, 2016; Lochmann & Deneve, 2011). For example, to cast BP into PLIF amounts to restricting BP’s message-passing within submodels of the kind suggested by Lemma 1. In other words, assuming that BP is to be adopted as the inference scheme, upon being presented with a causal query, an IT according to Lemma 1 will be selected—at the *meta-level*—by the reasoner and the corresponding submodel, as suggested by Lemma 1, will be retrieved, over which inference will be carried out using BP. This will lead to obtaining lower and upper bounds on the query, as reported in Lemma 1. If time permits, the reasoner builds up incrementally on the thus-far retrieved submodel so as to obtain tighter bounds on the query.¹² MCMC-based inference methods can be cast into PLIF in a similar fashion.

The problem of what parts of a CBN are relevant and what are irrelevant for a given query, according to (Geiger, Verma, & Pearl, 1989), was first addressed by Shachter (1988). The approaches proposed for identifying the relevant submodel for a given query fall into two broad categories (cf. Mahoney & Laskey, 1998, and references therein): (i) top-down approaches, and (ii) bottom-up approaches. Top-down approaches start with the full knowledge of the underlying CBN and, depending on the posed query, gradually *prune* the irrelevant parts of the CBN. In this respect, top-down approaches are inevitably from “god’s eye” point of view—a characteristic which undermines their cognitive-plausibility. Bottom-

up approaches, on the other hand, incrementally construct a submodel (by moving backwards from the query variables), using which the posed query can be computed. It is crucial to note that bottom-up approaches cannot stop at *intermediate* steps during the backward move and run inference on the thus-far constructed submodel without running the risk of compromising some of the (in)dependence relations structurally encoded in the CBN, which would yield erroneous inferences. This observation is due to the fact that there exists no local signal revealing how the thus-far retrieved nodes are positioned relative to each other and to the to-be-retrieved nodes—a shortcoming circumvented in the case of PLIF by introducing PL. It is worth reiterating again that PLIF subscribes to what we call the concurrent approach to reasoning (as opposed to the flawed sequential approach mentioned earlier), whereby retrieval and inference take place *in tandem*. The HMM example analyzed in Sec. 4.1, with infinitely large relevant submodel, stresses the importance and shows the efficacy of the concurrent approach.

Work on causal judgment provides support for the so-called alternative neglect, according to which subjects tend to neglect alternative causes to a much greater extent in predictive reasoning than in diagnostic reasoning (Fernbach & Rehder, 2013; Fernbach, Darlow, & Sloman, 2011). Alternative neglect, therefore, implies that subjects would tend to ignore parts of the relevant submodel while constructing it. Recent findings, however, seem to cast doubt on alternative neglect (Cummins, 2014; Meder, Mayrhofer, & Waldmann, 2014). Meder et al. (2014), Experiment 1 demonstrates that subjects appropriately take into account alternative causes in predictive reasoning. Also, Cummins (2014) substantiates a two-part explanation of alternative neglect according to which: (i) subjects interpret predictive queries as requests to estimate the probability of the effect when only the focal cause is present, an interpretation which renders alternative causes irrelevant, and (ii) the influence of inhibitory causes (i.e., disablers) on predictive judgment is underestimated, and this underestimation is incorrectly interpreted as neglecting of alternative causes. Cummins (2014), Experiment 2 shows that when predictive inference is queried in a manner that more accurately expresses the meaning of noisy-OR Bayes net (i.e., the normative model adopted by Fernbach et al. (2011)) likelihood estimates approached normative estimates. Cummins (2014), Experiment 4 shows that the impact of disablers on predictive judgments is far greater than that of alternative causes, while having little impact on diagnostic judgments. PLIF commits to the retrieval of enablers as well as disablers. As mentioned earlier, PLIF abstracts away from the inference scheme operating on the retrieved submodel, and, hence, leaves it to the inference scheme to decide how the retrieved enablers and disablers should be weighted and subsequently integrated. In this light, PLIF is consistent with the results of Experiment 4 in (Cummins, 2014).

In an attempt to explain violations of screening-off reported in the literature, Park and Sloman (2013) find strong

¹²The very property that the submodel gets constructed incrementally in a nested fashion guarantees that the obtained lower and upper bounds get tighter as the reasoner adopts smaller ITs; see Fig. 3(left).

support for the contradiction hypothesis followed by the mediating mechanism hypothesis, and finally conclude that people do conform to screening-off once the causal structure they are using is correctly specified. PLIF is consistent with these accounts, as it adheres to the assumption that reasoners carry out inference on their *internal* causal model (including all possible mediating variables and disablers), not the potentially incomplete one presented in the cover story; see also (Rehder & Waldmann, 2017; Sloman & Lagnado, 2015).

Experiment 5 in (Cummins, 2014), consistent with (Fernbach & Rehder, 2013), shows that causal judgments are strongly influenced by memory retrieval/activation processes, and that both number of disablers and order of disabler retrieval matter in causal judgments. These findings suggest that the CFP and memory retrieval/activation are intimately linked. In that light, next, we intend to elaborate on the rationale behind adopting the term “retrieve” and using it interchangeably with the term “consult” throughout the paper; this is where we relate PLIF to the concepts of Long Term Memory (LTM) and Working Memory (WM) in psychology and neurophysiology. Next, we elaborate on how PLIF could be interpreted through the lenses of two influential models of WM, namely, Baddeley and Hitch’s (1974) Multi-component model of WM (M-WM) and Ericsson and Kintsch’s (1995) Long-term Working Memory (LTWM) model. The M-WM postulates that “long-term information is downloaded into a separate temporary store, rather than simply activated in LTM,” a mechanism which permits WM to “manipulate and create new representations, rather than simply activating old memories” (Baddeley, 2003). Interpreting PLIF through the lens of the M-WM model amounts to the value for IT being chosen (and, if time permits, updated so as to obtain tighter bounds) by the central executive in the M-WM and the submodel being incrementally “retrieved” from LTM into M-WM’s episodic buffer. Interpreting PLIF through the lens of the LTWM model amounts to having no retrieval from LTM into WM and the submodel suggested by Lemma 1 being merely “activated in LTM” and, in that sense, being simply “consulted” in LTM. In sum, PLIF is compatible with both of the narratives provided by the M-WM and LTWM models.

A number of predictions follow from PL and PLIF. For instance, PLIF makes the following prediction: Prompted with a predictive or a diagnostic query (i.e., $\mathbb{P}(e|c)$ and $\mathbb{P}(c|e)$, respectively), subjects should not retrieve any of the effects of *e*. Introspectively, this prediction seems plausible, and can be tested, using a similar approach to (Cummins, 2014; De Neys, Schaeken, & d’Ydewalle, 2003), by asking subjects to “think aloud” while engaged in predictive or diagnostic reasoning. Also, PL yields the following prediction: Upon intervening on cause *c*, subjects should be sensitive to *when* effect *e* will occur, even in settings where they are not particularly instructed to attend to such temporal patterns. Recent findings suggesting that people have expectations about the delay length between cause and effect already provide some supporting evidence for this prediction (Greville & Buehner,

2010; Buehner & May, 2004).

There is a growing acknowledgment in the literature that, not only time and causality are intimately linked, but that they *mutually constrain* each other in human cognition; cf. (Buehner, 2014). In line with this view, we see our work also as an attempt to formally articulate how time could guide and constrain causal reasoning. While many questions remain open, we hope to have made some progress towards better understanding of the CFP at the algorithmic level.

Acknowledgments. We are grateful to Thomas Icard for valuable discussions. We would also like to thank Marcel Montrey and Peter Helfer for helpful comments on an earlier draft of this work. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada under grant RGPIN 262017.

References

- Baddeley, A. (2003). Working memory: looking back and looking forward. *Nature Review Neuroscience*, 4(10), 829–839.
- Baddeley, A. D., & Hitch, G. (1974). Working memory. *Psychology of Learning and Motivation*, 8, 47–89.
- Bramley, N. R., Gerstenberg, T., & Lagnado, D. A. (2014). The order of things: Inferring causal structure from temporal patterns. In *Proceedings of the 36th annual conference of the cognitive science society* (pp. 236–241).
- Buehner, M. J. (2014). Time and causality: editorial. *Frontiers in Psychology*, 5, 228.
- Buehner, M. J., & May, J. (2004). Abolishing the effect of reinforcement delay on human causal learning. *Quarterly Journal of Experimental Psychology Section B*, 57(2), 179–191.
- Cummins, D. D. (2014). The impact of disablers on predictive inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(6), 1638.
- De Neys, W., Schaeken, W., & d’Ydewalle, G. (2003). Inference suppression and semantic memory retrieval: Every counterexample counts. *Memory & Cognition*, 31(4), 581–595.
- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102(2), 211.
- Fernbach, P. M., Darlow, A., & Sloman, S. A. (2011). Asymmetries in predictive and diagnostic reasoning. *Journal of Experimental Psychology: General*, 140(2), 168–185.
- Fernbach, P. M., & Rehder, B. (2013). Cognitive shortcuts in causal inference. *Argument & Computation*, 4(1), 64–88.
- Fodor, J. A. (1987). Modules, frames, fridges, sleeping dogs, and the music of the spheres.
- Geiger, D., Verma, T., & Pearl, J. (1989). d-separation: from theorems to algorithms. *Fifth Workshop on Uncertainty in Artificial Intelligence*, pp. 118–125.
- Gershman, S. J., & Beck, J. M. (2016). Complex probabilistic inference: From cognition to neural computation. In *Computational Models of Brain and Behavior*, ed A. Moustafa.
- Glymour, C. (1987). Android epistemology and the frame problem, *The Robot’s Dilemma: The Frame Problem in AI*, pp. 65–75.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: causal maps and bayes nets. *Psychological Review*, 111(1), 3–32.
- Greville, W. J., & Buehner, M. J. (2010). Temporal predictability facilitates causal learning. *Journal of Experimental Psychology: General*, 139(4), 756.
- Hume, D. (1748/1975). *An inquiry concerning human understanding*. Oxford University Press.
- Icard, T. F., & Goodman, N. D. (2015). A resource-rational approach to the causal frame problem. *Proc. of the 37th Annual Meeting of the Cognitive Science Society*.
- Lagnado, D. A., & Sloman, S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(3), 451–60.
- Lochmann, T., & Deneve, S. (2011). Neural processing as causal inference. *Current Opinion in Neurobiology*, 21(5), 774–781.
- Mahoney, S. M., & Laskey, K. B. (1998). Constructing situation specific belief networks. *Proc. of the Conference on Uncertainty in Artificial Intelligence*, pp. 370–378.
- Marr, D. (1982). *Vision: a computational approach*.
- Meder, B., Mayrhofer, R., & Waldmann, M. R. (2014). Structure induction in diagnostic causal reasoning. *Psychological Review*, 121(3), 277.
- Park, J., & Sloman, S. A. (2013). Mechanistic beliefs determine adherence to the markov property in causal reasoning. *Cognitive Psychology*, 67(4), 186–216.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann.
- Rehder, B., & Waldmann, M. R. (2017). Failures of explaining away and screening off in described versus experienced causal learning scenarios. *Memory & Cognition*, 45, 245–260.
- Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in Cognitive Sciences*, 20(12), 883–893.
- Schlottmann, A. (1999). Seeing it happen and knowing how it works: how children understand the relation between perceptual causality and underlying mechanism. *Dev Psychol*, 35(1), 303.
- Shachter, R. D. (1988). Probabilistic inference and influence diagrams. *Operations Research*, 36(4), 589–604.
- Simon, H. A. (1957). *Models of man*. Wiley.
- Sloman, S. A., & Lagnado, D. (2015). Causality in thought. *Annual Review of Psychology*, 66, 223–247.