

To do or not to do: finding causal relations in smart homes

Kanvaly Fadiga^{*†}, Étienne Houzé^{*‡}, Ada Diaconescu^{*}, Jean-Louis Dessalles^{*}

^{*}LTCI Lab, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France

Email: {first name}.{last name}@telecom-paris.fr

[†]Ecole Polytechnique, Institut Polytechnique de Paris, Palaiseau, France

Email: {first name}.{last name}@polytechnique.edu

[‡]EDF R& D, Palaiseau, France

Email: {first name}.{last name}@edf.fr

Abstract—Research in Cognitive Science suggests that humans understand and represent knowledge of the world through causal relationships. In addition to observations, they can rely on experimenting and counterfactual reasoning – i.e. referring to an alternative course of events – to identify causal relations and explain atypical situations. Different instances of control systems, such as smart homes, would benefit from having a similar causal model, as it would help the user understand the logic of the system and better react when needed. However, while data-driven methods achieve high levels of correlation detection, they mainly fall short of finding causal relations, notably being limited to observations only. Notably, they struggle to identify the cause from the effect when detecting a correlation between two variables. This paper introduces a new way to learn causal models from a mixture of experiments on the environment and observational data. The core of our method is the use of selected interventions, especially our learning takes into account the variables where it is impossible to intervene, unlike other approaches. The causal model we obtain is then used to generate Causal Bayesian Networks, which can be later used to perform diagnostic and predictive inference. We use our method on a smart home simulation, a use case where knowing causal relations pave the way towards explainable systems. Our algorithm succeeds in generating a Causal Bayesian Network close to the simulation’s ground truth causal interactions, showing encouraging prospects for application in real-life systems.

Index Terms—Causal Structure Discovery, Smart Home, Causal Inference

I. INTRODUCTION

Self-Adaptive Systems (SAS) are by nature interacting with a changing environment, be it software or physical [29]. In this context of interactions, the ability to model the environment is prime, as it would help to trace back failures, identifying conflicts between goals, or perform explanatory reasoning. [13] has shown that, in typical smart home setups, explaining decisions to the user reduces the risk of wrong interventions. However, identifying causal relations in the environment of a SAS is a hard task.

Hard-coding the causal model, i.e. expressing constraints and links upon variables as rules or a static ontology is possible, but shows limited interest in the case of SAS. Indeed, since adaptability to a changing environment is a core feature of SAS, a static model of the environment is not suited to this configuration. Operating changes to the model could be

considered, but is likely to require many human interventions, thus contradicting the principles of autonomic computing [10].

To illustrate this issue, consider the following situation. A user is experiencing an anomaly in the temperature control system of her smart home, as the temperature is unexpectedly cold. A hard-coded model has been implemented, which contains causal links from heater or thermometer malfunctions to the mishandling of the temperature. However, both these possibilities are discarded, as no component seems to report any problem. In this case, the cause might lie in an unexpected relation: as the user recently moved the lamp closer to the temperature sensor, and that the days, in winter, are shorter, the light is on, which produces heat, effectively skewing thermometer measures. This configuration being particular to this home, no hard-coded prior causal knowledge would be able to anticipate it without ad-hoc rules.

To overcome this common pitfall, many recent smart home systems integrate Machine Learning components to predict the environment’s behavior and make optimized decisions [11]. However, spurious correlations are often found in data, especially in high dimensions [3], leading to misinterpretation and erroneous causal relations. These approaches thus mostly fall short of providing the user with a comprehensible causal model of the environment.

The theory of Causality, brought to attention among others by J. Pearl [20, 18], offers tools to identify and eliminate spurious correlations in the construction of a causal model, mostly by formalizing the concept of *intervention* defined by “do-operation”. Our method consists of augmenting a standard Machine Learning approach with interventions on selected variables to infer causal relations. These interventions are carried out with the aim of conducting a counterfactual experiment by observing what happens when we do and do not do an action, hence the title: To do or not to do. The result is a Causal Bayesian Network, i.e. a Bayesian Network whose structure is a causal graph of the environment.

Our approach’s advantage is that it considers that interventions are not always possible. In ideal setups, where all variables are operable, it yields results close to ground truth causal relations. It is still able to perform well when fewer variables are available for intervention, being closer to the

performance of data-driven methods.

The algorithm we will present is generic and can be applied to build causal models of various environments. But it can be computationally expensive to apply to an environment with a very large network. That is why we choose to apply it in the smart home case because it offers many advantages at computation level. Firstly, we don't start from scratch as we can begin with a hard-coded causal model and then incrementally improve it. Moreover, making interventions in a smart home is easier than in some environments (e.g. a nuclear power plant). Furthermore, the area of influence of some variables is limited to their neighborhood, which reduces the number of relationships to consider. Smart homes also provide good examples of closed environments managed by SAS. As such, they also provide simulators, which can be used to implement an intervention operator without being limited by common physical constraints (time, safety issues, incompatibilities). In addition, they can present unusual or surprising situations where the use of a causal diagnostic can help intervene on the system to improve performance [13, 7].

The rest of this paper is organized as follows. In section II, we present the theoretical bases of causality and Bayesian graphs upon which our approach is built. Then, we review some existing approaches to related issues in section III. We then detail our method, and propose, in section IV to illustrate it by comparing known causal graphs of a smart home with the results of our method in section V. Finally, we analyze the current limits of our approach and see how it can be integrated into broader systems in section VI.

Level	Activity	Question	Example
Association $P(Y X)$	Seeing	How would seeing X change my belief in Y ?	What is the probability of someone's presence knowing the light is on ?
Intervention $P(Y do(X))$	Doing	What if I do $X = x$?	If I turn the heater on, would the temperature change?
Counterfactual $P(Y_x X', Y')$	Imagining	What if things were different?	Had the light been off, would the temperature had changed?

TABLE I: Pearl causality hierarchy [20, 1]

II. THEORETICAL BACKGROUND

Many examples from Machine Learning point out that algorithms usually lack the understanding of causal relations behind observations and predictions, causing misinterpretations of correlations [17]. [20] goes further by integrating this observation into a "ladder of causation", in which three distinct levels are identified (Fig 1):

- **Observing** corresponds, according to Pearl [20], to the first and more reachable level of cognition: observing the world and noting correlations, dependence between some sets of variables. This stage is the ground for many modern AI approaches based on data analysis.
- **Acting** This advanced stage of cognition requires the agent to be able to act on some variables of the environment, observe the consequences and infer causal

relations. The typical question answered at this point is "If I do this, what will happen next ?"

- **Counterfactual Thinking** involves consideration of an alternate version of a past event, or what would happen under different circumstances for the same experiment. According to Pearl, this level of cognitive ability is only reached by humans. The typical question would be : "What if the apple was twice as heavy ? Would it have fallen at a different speed?"

During the twentieth century, from the causal chains of Wright [30] to the integration of causal inference into Machine Learning algorithms [21, 22], research in the field of Causality Theory aimed to formalize the intuitive concepts of cause-consequence relations.

A. Structural Causal Model

Causal models aim at representing the interactions between cause and consequence without ambiguity. From the definition of [21], a Structural Causal Model (SCM) contains $C \rightarrow E$ if and only if $C \equiv N_C$ and $E \equiv f_E(C, N_E)$. That is if the cause variable C can be assigned to some specific random distribution N_C while E can be computed from a deterministic function of C and some random noise N_E . Note that, in this definition, both the causal relation f_E and the effect noise N_E are independent from the cause C .

For more complex systems, where causal dependencies between variables may be multiple, we can use a *causal diagram* to represent an underlying structural causal model. A causal diagram (Fig 1) is a directed acyclic graph [20, 21, 25] that shows the causal relationships between variables. The nodes of the graph are the variables, and an edge (C, E) belongs to the graph if and only if $C \rightarrow E$ belongs to the underlying SCM. For example, in the diagram presented in Fig 1, the arrow connecting variables *Heater* and *Temperature* ($H \rightarrow T$) indicates that the temperature is causally influenced by the state of the heater. Another point of view is to consider *Temperature* as listening to the *Heater* variable to choose its value.

Compared to the more general definition of SCM, causal diagrams add the condition of being acyclic [18], encompassing the idea that causality flows in one direction only: if C has a causal influence on E , then E cannot influence C . This further prevents a variable from influence itself.




Fig. 1: A simple causal diagram.

B. Do-operator

The idea of being able to intervene in the environment to test a causal relationship between variables is prime in the

literature of Causality Theory [18, 21] and can be linked to a general controlled environment experiment.

The intervention operation has been formalized by Pearl by introducing the *do-calculus* [18, 20]. Following his notations, $do(C = x)$ means that C has been forced to take the value x by an external action. It follows that, if $C \rightarrow E$ was part of the SCM, the causal relation $E = f_E(C, N_E)$ remains unchanged by this operation. This operation therefore allows to identify causal relations: if $\mathbb{P}(E) \neq \mathbb{P}(E|do(C = x))$, there is a causal connection $C \rightarrow E$. In this case, we will use the notation $do(C) \rightsquigarrow E$.

While mere observations of the variables H, T and W from the example of Fig 1 would show correlations between H and T , interventions would give more details on the underlying SCM. On one hand, $\mathbb{P}(Heater | do(Temperature = 20)) = \mathbb{P}(Heater)$ and on the other hand, $\mathbb{P}(Temperature | do(Heater = High)) \neq \mathbb{P}(Temperature)$, reflecting that the heater causes the temperature change, not the other way around.

The do-operation $do(C = x)$ considers an *external* intervention, meaning that it forces the variable C to a given value x , while making it insensitive to all other variables. On a causal diagram, this is equivalent to removing all incoming edges to node C . For instance, if we consider the simple causal diagram $C_0 \rightarrow C_1 \rightarrow E$, performing the intervention $do(C_1 = x)$ will remove the edge $C_0 \rightarrow C_1$, thus making both C_1 and E independent from C_0 , thus revealing the linear structure of the graph.

C. Bayesian Network

As Causality Theory emerged with causal diagrams, links can be made with *Bayesian Networks* which are a broadly used tool for representing and modeling correlated variables [9]. Numerous methods for training and dynamically building Bayesian Networks in many different application contexts exist in the literature [2, 9].

Formally, a Bayesian Network (see Fig 2) is a directed acyclic graph (DAG) where the nodes correspond to random variables. Each node is associated with a set conditional probabilities $\mathbb{P}(X_i | par(X_i))$, where X_i is the variable associated with the specific node and $par(X_i)$ denotes the set of parents of node X_i .

To build a Bayesian network, one needs to:

- define the graph of the model, i.e. the different variables, and their links,
- find, for each of these variables, the table of probabilities conditioned on its parent variables

The graph is also called the "structure" of the model, and the probability tables its "parameters". Structure and parameters can be provided by experts, or calculated from data, but in general the structure is defined by experts and the parameters calculated from experimental data.

A Bayesian network carries no assumption that the arrow has any causal meaning. However, Bayesian networks hold the key that enables causal diagrams to interface with data. The probabilistic properties of Bayesian networks and the

belief propagation algorithms that we will present later are indispensable for understanding causal inference.

The main differences between Bayesian networks and causal diagrams lie in how they are constructed and the uses to which they are put. A Bayesian network is literally nothing more than a compact representation of a huge probability table. The arrows mean only that the probabilities of child nodes are related to the values of parent nodes by a certain formula (the conditional probability tables) and that this relation is sufficient. That is, knowing additional ancestors of the child will not change the formula. Likewise, a missing arrow between any two nodes means that they are independent, once we know the values of their parents.

If, however, the same diagram has been constructed as a causal diagram, then both the thinking that goes into the construction and the interpretation of the final diagram change. In the construction phase, we need to examine each variable, say C , and ask ourselves which other variables it "listens" to before choosing its value. The chain structure $A \rightarrow B \rightarrow C$ means that B listens to A only, C listens to B only, and A listens to no one; that is, it is determined by external forces that are not part of the model.

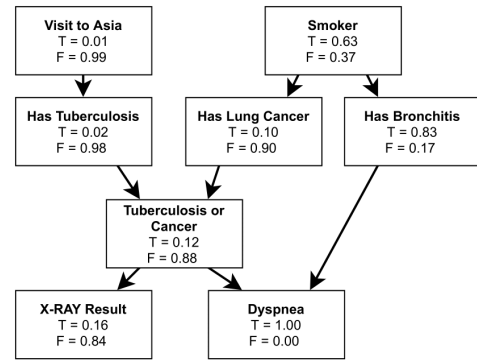


Fig. 2: A simple Bayesian network, known as the Asia network¹

III. RELATED WORK

Over recent years, different approaches have tried to close the gap between "classical" observation-based Machine Learning and Causal Theory. For instance, Reinforcement Learning, as already noted by Pearl [20] can be seen as a better approach than pure correlation observations, since the agent has the opportunity to act on its environment and learn from its reactions [26]. Thus, Reinforcement Learning has proved very powerful in tasks that were previously considered as requiring intelligent thinking, such as games [24]. An agent in a reinforcement learning environment provides intervention data for learning a causal model through the exploration of the state space. However, the causal model it learns is implicit therefore it cannot be used and interpreted.

[16] uses a different approach to learning Bayesian Networks, trying to identify a minimal equivalence class between DAGs that fit with the observation data. The result is then

¹From *Bayes Server* (bayesserver.com/examples/networks/asia)

presented as a Partially Directed Acyclic Graph (PDAG) which is a directed graph with some undirected connections. While this method offers the advantage of keeping the graph simple and shows good predictive performance, it still relies on mere observations, and as such lacks causal information that may impact its interpretation. [15] completes the previous method by orienting the undirected connections of the PDAG through interventions. He ends up orienting relations proposed by the correlations which are often non-causal relationships. Moreover, in their approach they did not take into account the fact that on some variables it is not possible to make interventions.

Some applications consider counterfactual reasoning and integrate it into the learning process of a SCM [14, 27]. In their workflow, they consider an agent that learns a causal model of its environment and then use this knowledge to perform counterfactual reasoning and improve performance. Results in providing explanations for an agent’s behavior in the controlled environment of a strategy game are encouraging [14]. In a closer-to-life situation, [22] found that allowing do-operations in a learning framework could improve performance in a classification task and achieve better-than-humans detection of a medical condition.

Our approach differs from [15, 16], as it accounts for some variables being unavailable to intervention. Instead, we learn causal relations using a mix of observation data and, when possible and useful, intervention data.

IV. LEARNING CAUSAL BAYESIAN NETWORKS WITH INTERVENTION

The base intuition for our approach is to test whether intervention on one variable C has an influence over another variable E , observed as a change in their distribution. If so, we know from Causality Theory that there is a causal relation $C \rightarrow E$ in the SCM of the system, but an ambiguity remains whether this relation is direct or not. We propose to incrementally block causal paths of nodes connected to the node on which we act, effectively narrowing down the possible relations.

We illustrate our approach in a setup consisting of Boolean variables. For illustration purposes, let’s consider the simple situation displayed in Fig 3: a room whose temperature (hot or cold) is influenced by the state of a heater (on or off). The heater can be triggered by the user’s presence in the room. Similarly, the window can be either open or close.

A. Testing causal influence using interventions

1) *Direct Influence:* Testing the direct influence between two variables C and E is answering the following question: “Does C have an influence on B ?”. Our approach to this question is to incrementally remove possible causal relations following different interventions. These interventions are made by directly acting upon the environment and monitoring possibly influenced variables for changes in their probability distribution. To test possible changes, we use a chi-squared χ^2 test on the distributions $\mathbb{P}(E | do(C = 0))$ and $\mathbb{P}(E | do(C = 1))$.

This test allows removing non-causal connections between pairs of variables, using both intervention operation and counterfactual reasoning. The intervention operations can be performed by directly letting our algorithm act on selected variables in the environment, thanks to the preconditions we applied to the setup. For instance, in the example of Fig 3, the distribution of L changes depending on whether the person is detected inside ($P=1$) the room or not ($P=0$). Conversely, the distribution of P is not affected by the value assigned to L during the intervention operation. These two operations, therefore, lead to the conclusion that $do(P) \rightsquigarrow L$ is true and $do(L) \rightsquigarrow P$ is false.




Fig. 3: A room causal diagram




Fig. 4: Different intervention tests. (4a), the probability density distribution of L changes depending on whether the intervention sets P to 0(blue) or 1(orange). In (4b), interventions on L do not affect the probability distribution of P . (4c): intervening on P shows a change in the distribution of T . (4c): intervening on P shows a change in the distribution of T . However, conditioning this relation with $H = 0$ removes the relation(4d)

2) *Conditional Influence:* The case of evaluating a conditional causal influence can be summarized with the question: “Did C have an influence on E given B ?”($do(A) \rightsquigarrow C | do(B)$). As opposed to the previous case, the causal path is indirect and thus requires additional processing. Here, we process by testing if the causal influence between C and E still holds conditioned on the value of the third variable B . That is, checking if, for some value u taken by B , we would observe, via a chi-squared test, a difference between $\mathbb{P}(E | do(C = 0), do(B = u))$ and $\mathbb{P}(E | do(C = 1), do(B = u))$.




Fig. 5: Principle of our algorithm. (5a) ground truth causal model of the environment. (5b) Initialization to a fully-connected graph over the variables. Non-doable arrows and nodes are shown in blue. (5c) Causality tests with interventions remove the red arrows. (5d) Arrows are removed, either by independence test (in yellow) or causality test (in red). (5e) The final graph is obtained by removing cycles, prioritizing non-doable arrows.

This operation can be viewed as “locking” the value of B to a given value u , and observe if the causal relation holds. In the examples of Fig 4c and 4d, we infer the causal relations:

- $do(P) \rightsquigarrow T \mid do(L = 0) : \text{True}$
- $do(P) \rightsquigarrow T \mid do(H = 0) : \text{False}$

B. Causal Learner Algorithm

Our algorithm, presented in alg.1, iterates over the previously described elementary steps to remove non-causal pairs of variables. To this end, we start by considering a fully connected graph over all the variables in the system (see Fig 5). Then, selected causal influence tests are used to remove arrows for unrelated variables. These tests are performed by increasing order of conditioning: this allows to test the costlier high-order conditioned influences on few arrows, as many have already been discarded by the first series of tests.

As shown in Fig 5, a major limitation of this approach is that some do-operations are not feasible in realistic setups: in our example, this is the case for the temperature variable T , as one does not arbitrarily set the temperature of a room to some fixed value without modifying other variables (e.g. the heater state H). We therefore call the corresponding temperature node a *non-doable node*, and consider all of the outgoing relations as “*non-doable arrows*”, or ND-arrows, in the graph. These arrows are not directly removable since the corresponding do-operations cannot be performed.

Processing ND-arrows, therefore, requires another approach. First, similarly to the PC-algorithm from [25], we use a simple chi-squared test to identify whether the two variables are correlated since a causal relationship implies a correlation between variables. This first step allows to remove some connections, but, for the remaining connections, it does not provide any direction for the relation. Furthermore, one needs to be cautious about the potential risk of mislabeling correlations as causal relations. As such, remaining ND-arrows should be considered only as candidate causal relations.

To handle the rest of the process, we rely on the fact that the causal diagram is, by definition, a DAG. This condition leads to the removal of some arrows among the remaining candidates. Depending on the configuration, different possibilities are considered, as Fig II shows:

Algorithm 1 Extended *do* Causal Learning Algorithm

```

1: Initialization:
2:  $\mathcal{G}$  is the fully connected graph over nodes of  $\mathcal{X}$ 
3:  $k \leftarrow 0$ 
4: while There are nodes with more than  $k$  neighbors in  $\mathcal{G}$  do
5:   for each such node  $X_A$ , each of its neighbors  $X_B$  do
6:     for each subset  $\mathcal{S}$  of  $k$  neighbors of  $X_A$  do
7:       # Influence test for doable node
8:       if  $X_A$  is doable then
9:         Compute  $do(X_A) \rightsquigarrow X_B \mid do(\mathcal{S})$ 
10:        Remove  $A \rightarrow B$  from  $\mathcal{G}$  if need to be
11:       # Independence test for non-doable node
12:       else
13:         Compute  $Corr(X_A, X_B \mid \mathcal{S})$ 
14:         Remove  $X_A \rightarrow X_B$  if variables are independent
15:       end if
16:     end for
17:   end for
18:    $k \leftarrow k + 1$ 
19: end while
20: Postprocess  $\mathcal{G}$  to turn it into a DAG by removing least significant arrows.

```

	Present relationships	Potential output
1	<p>Causal arrow, A is the cause of B</p> <p>Exactly one of the following holds:</p> <ol style="list-style-type: none"> 1. A is the cause of B 2. B is the cause of A 3. No connection, there is a spurious correlation create by unmeasured confunder 	<p>undirected connection</p>
3	<p>A is the cause of B</p> <p>A is not the cause of B.</p> <p>Exactly one of the following holds:</p> <ol style="list-style-type: none"> 1. B might be the cause of A 2. No connection, there is a spurious correlation create by unmeasured confunder 	<p>Flagged connection</p>
5	<p>The confounder is known, so we can delete the spurious correlation created by C.</p>	

TABLE II: The different possible configurations for processing the remaining ND-arrows. Causal arrow are shown in green, ND-arrow are in blue. Red arrow represent arrow remove by causality test

- **case 1:** As no ambiguity exists, the arrow is kept in the graph.
- **case 2:** No information can be gathered through correlation study alone. If no direction creates a cycle in the graph, the algorithm will keep the undirected relation,

and tag it as potentially spurious. Further data may lead to eliminating both of the arrows.

- **case 3:** One direction of the relation has been tested through a do-operation, while the other has not. The algorithm will therefore keep the direction that has been tested with an intervention.
- **case 4:** While this ND-arrow can be a spurious correlation, the algorithm will keep it if it does not create a cycle in the resulting graph. It will however be flagged as such. Otherwise, the arrow is removed from the graph. More generally, if keeping several ND-arrows would lead to a cycle, the algorithm will remove the least significant one with respect to Chi-square score.
- **case 5:** In this case we see an ND-arrow that can be preserved if there is a confounder that creates a correlation between A and B. Here we see that C is a confounder. So we drop the ND-arrow.

After processing all ambiguous cases, the algorithm outputs a DAG representing a causal model compatible with observations from the system. This causal diagram can then be used as a basis for further analysis. A first possibility is to use it to infer potential causes to unusual situations. As such, it may be included in a broader-scoped explanation process [7]. A second prospect, detailed here, is to use this diagram as the structural basis of a Bayesian Network for finer causal inference.

Our algorithm requires an exponential number of tests with respect to the number of variables (node). But, as mentioned in I, in the case of the smart home, we can consider an incremental approach. A base, generic, causal model is implemented, which is later refined by testing its connections with our method. Then, as new devices are added to the system, we consider only the newly created interactions, thus incrementally augmenting the causal graph.

C. From Causal Model to Bayesian Network

In the literature, training a Bayesian is usually divided into two main parts [9]: learning the structure of the graph and estimating its parameters. Since we use the previously learned causal diagram as a base structure, we will only focus in this part on learning the different parameters of the network, i.e. the probability tables for each node conditioned on its parents. We will call the resulting graph a *Bayesian Causal Network* to emphasize its particular structure: while usual Bayesian Networks do not entail causality between their nodes, our approach leads to a graph whose connections entail a cause-effect relation.

To estimate these parameters, a conventional approach is to use a maximum likelihood estimator [4], which can be limited to estimating variables values given their parents' values from past observational data only. For example, if we consider the graph from Fig 6, we would compute p_{00} with:

$$p_{00} = \frac{N_{T=0,(W,H)=(0,0)}}{N_{T=0,(W,H)=(0,0)} + N_{T=1,(W,H)=(0,0)}} \quad (1)$$

where $N_{T=i,(W,H)=(j,k)}$ is the number of past occurrences of $(T, W, H) = (i, j, k)$.




Fig. 6: Example of a small Bayesian network. The probability table for node T is displayed.

However, this conventional approach might be facing issues for some estimations, notably if the number of occurrences is small. For instance, in our example of Fig 6, it might not be possible to estimate the parameter p_{01} . The introduction of the do-operation removes this limitation, as it becomes possible, for doable nodes, to generate all kinds of situations required to observe the outcome and estimate missing parameters of the Bayesian Network.

D. Causal Inference on Causal Bayesian Network

Upon completion of the training, the resulting Causal Bayesian Network can be seen as a “conditional probability machine” [9]. It can be used for different tasks requiring inferring new knowledge on the system. For instance, [31] shows how a Bayesian Network can be used to compute the probabilities of different activities compatible with the observed interaction of a user with IoT device in a smart home. This example shows the different possibilities offered by a Bayesian Network: diagnostic and predictive inference.

- **Predictive:** This kind of inference is interested in “guessing” the most probable future state of the system, given a configuration, i.e. answering the question: “*What happens if X is equal to x?*” As Fig 7a shows, if evidence is put on node P, the inference will propagate following the direction of causal arrows, to the children of the affected node P.
- **Diagnostic:** On the other hand, diagnostic inference is interested in looking into the probable causes of observed consequences: “*what would be the cause of X = x?*” The inference therefore goes backwards, as displayed in Fig 7b: from the observation on L, we infer the probable state of P, which will entails consequences over H and T.

In either case, inference works as follows: we denote by $Bel(X = x)$ the belief that a node takes a given value (see Fig 2, where beliefs are displayed for each node). Following observation of the system, the beliefs of one or several nodes are set to a set value. For instance, in Fig 6, knowing that the person is present will set the value of P to 1 with a probability 1. This change of beliefs will then be propagated through the graph, following Bayes' rule.

While we let the details of the propagation algorithm out of the scope of this paper (readers interested in a complete




Fig. 7: Belief propagation in a Bayesian Network can be either forwards (7a) for predictive applications, backwards (7b) for diagnostic purpose, or oriented both ways (7c).

description of the process may refer to [19, 9]), we could visualize the propagation mechanism as follows. The propagation algorithm is iterative. At every step, each node X passes the following messages: to its children Y , $\pi_X(Y)$ containing transition probabilities; to its parents Z , $\lambda_X(Z)$ containing likelihood information. Conversely, it receives messages $\pi_Z(X)$ from its parents, and $\lambda_Y(X)$ from its children (Fig7). Each node then updates its beliefs according to the messages it receives:

$$Bel(X) = \alpha \lambda(X) \pi(X) \quad (2)$$

where α is a normalizing factor, $\lambda(X) = \prod_Y \lambda_Y(X)$ and $\pi(X) = \prod_Y \pi_Y(X)$ are the products of all messages received from children and parents, respectively. As shown by [19], for DAGs, this propagation method converges to the belief values satisfying the observations and the network's parameters in a finite number of steps.

The predictive and diagnostic inference then allows answering various queries about the environment without having to further intervene on the system. Applications of such knowledge are further discussed in sec. VI. One might however note that, as opposed to a traditional Bayesian Network, our proposed CBN uses only causal relations. As such, one might argue that the entailed reasoning appears more “natural”, a case confirmed by the observation that causal relations are algorithmically simpler [21, 8].

node	($x=0$, $x=1$)
P	(0.5, 0.5)
L	(0.45, 0.55)
H	(0.7, 0.3)
T	(0.67, 0.33)
W	(0.7, 0.3)

TABLE III: Prior probability of each node of 7b

node	($x=0$, $x=1$)
P	(0.89, 0.11)
L	(1, 0)
H	(0.86, 0.14)
T	(0.73, 0.27)
W	(0.7, 0.3)

TABLE IV: Posterior probability after $L=0$.

V. EXPERIMENTS AND RESULTS

A. General Workflow

The following section describes in detail our workflow in a smart home simulator, from data generation to training the CBN and using it for inference tasks.

1) *Smart Home simulator*: Our experiments are built upon the iCasa platform [12]. Based on the OSGi framework, it offers a service-oriented platform for simulations of smart home physical systems. Its autonomic manager keeps track of

Boolean variable	Simulation measure
person P	(User.x, User.y) \in room
outdoor O	outdoor.temperature \geq threshold
light L	light.powerStatus = 1
presence Pr	sensor.presenceSensed = 1
power Pow	house.powerConsumption \geq threshold
thermometer T	thermometer.temperature \geq threshold
window W	window.open = 1

TABLE V: Correspondence between simulation measures and the Boolean variables.

currently used devices, which allows for runtime deployment and modification of configuration. This enables the simulation of scenarios where variable interactions are more intricate. In our example, we simulated the behavior of different rooms, each one characterized by physical variables such as temperature and illumination. Each of these rooms is equipped with different devices able to monitor or modify the room's physical variables: heater, thermometer, presence sensor, light, etc. Table V shows the different monitored variables of the example. The entire configuration is shown in Fig 8 using the iCasa Web UI.

Using a simulation, as opposed to using a real setup, brings two main advantages for our experiments. First, it allows having a perfect knowledge of the ground-truth causal interactions, as they are directly encoded into the simulator. Secondly, it provides easy control over different parameters and thus allows to perform, if desired, some interventions that would not be feasible in real life. This will allow us to test the effect of having access to more or less possible interventions for our algorithm.



Fig. 8: The iCasa GUI showing the basic setup for our experiments: four rooms equipped with a presence sensor triggering heating and lighting, and a thermometer.

2) *Observation Data Generation*: Once the initial setup is complete, we let the simulation run while the various devices are left in “autonomous mode”, i.e. they are able to adapt to changes to maintain some key environment variables within a target range, for instance, the temperature and CO2 concentration of each room. At runtime, we randomly act on some of these variables or components to observe how the system reacts to change. In total, our continuous observation generated around 500 data points. Since values from variables are originally numerical, we convert them to Boolean values

by using simple threshold comparisons. Thus, we obtain a set of Boolean observations which are used to observe correlations between variables.

3) *Intervention Data Generation*: To perform interventions, we use the possibility offered by the simulator setup to disable some devices. Disabled devices will no longer react to their input sensors, thus achieving the Markov blanket independence implied by intervention [18]. Then the value of the device is set to a fixed value. For instance, the intervention $do(Heater = 1)$ will cause the heater to turn on while being insensitive to any environmental factor such as the detection of the user’s presence.

To generate intervention data, we then proceed as follows: we sample the house in a state s where each variable is assigned a value, and from this state, make an intervention $do(X = x)$ on a selected variable. We did 20 interventions per node at each stage. After a set time period Δ_t , we measure the resulting state of the house s' , considering only variables of interest (variables correlated to X). The period Δ_t is set to a given value manually chosen from prior experiments with the simulator, to allow the system to reach an equilibrium state after the intervention. We will discuss further time considerations in sec. VI.

B. Results

The causal model of the simulation we used for our experiments is displayed in Fig 9a. We first consider three variables, namely the presence sensor, the house’s power consumption, and the room’s temperature as ND-nodes. While the simulation setup would allow us to intervene on them, we introduced this limitation to observe the impact of ND-nodes on causal discovery.

1) *Causal Bayesian Network learning*: The construction of the causal graph is shown in Fig 9. First, observations of correlated variables and results of interventions yields a “raw” output depicted in fig 9b. Note how the presence of ND-nodes introduces ND-arrows emerging from them. The next step of our algorithm processes this raw output to remove the least significant arrows to make it a DAG that is compatible with the observations. The output of this step, shown in Fig 9c, contains two arrows flagged as ND-arrows. When comparing this final output to ground truth, in Fig 9d, we notice that one of these ND-arrows was erroneous, displaying a performance limit in the case of ND-nodes. On the other hand, one causal relation, between light status and power consumption, was missed by our approach. This mistake can be explained in this situation, by the relatively small impact of light, in comparison to the heater, on power consumption.

Building on the structure of the causal graph presented in Fig 9, we complete the learning process by using maximum likelihood estimates to finally provide a Causal Bayesian Network.

2) *Impact of ND-nodes*: To analyze the impact of ND-nodes on the learning process, we have explored different arrangements of ND-nodes, by allowing or not variable modifications in the simulation. Figure 10 shows the resulting graphs

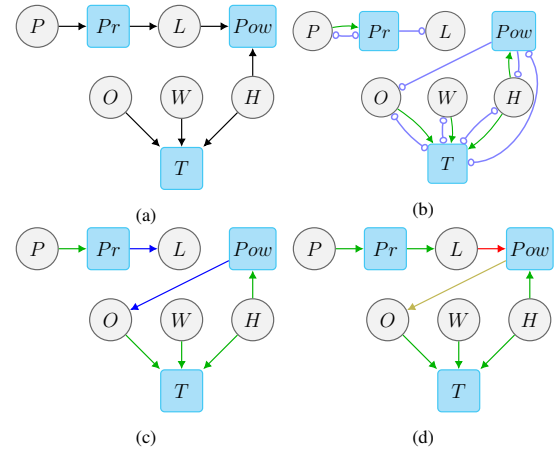


Fig. 9: Results of our approach applied to the smart home simulation. Fig 9a shows Groundtruth causal model for the living room. Non-doable variables are shown in blue. The raw output of conditional testings, shown in (9b), is then processed to remove less significant arrows to obtain a DAG (9c). (9d): comparison between this output and the ground truth diagram from Fig 9a: the red arrow is a missed relation while the yellow one is a connection wrongly added to the model.

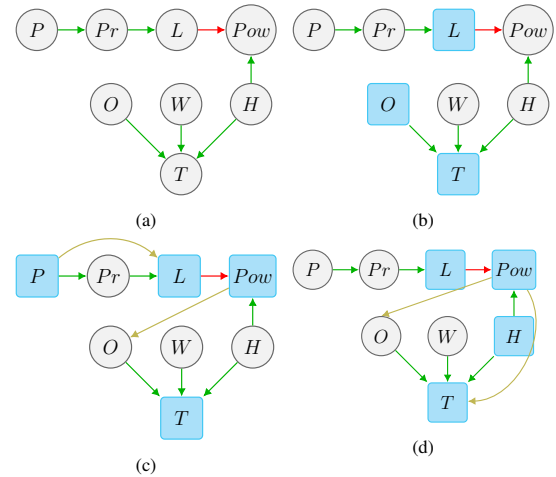


Fig. 10: Causal graphs obtained with different configuration of non-doable nodes (in blue). Comparison with the ground-truth relation (9a: green arrows are correct, while red, yellow and blue respectively denote missed, wrongly added or bidirectional connections).

our algorithm learns, depending on the do-ability of nodes. Overall, relations between non-doable nodes are affected, while the algorithm stays robust concerning the other relations in the graph.

3) *Unusual Causal Relations*: A motivation for learning CBN with minimal prior knowledge was the ability to adapt to unusual situations, detect them and use knowledge to provide explanation to the user. Such a situation may occur in our experimental setup: the thermometer has been implemented to be sensitive to the heat produced by nearby devices (notably the light or the heater). In our setup, four rooms are simulated with the same devices, however the precise location of devices within each room is random. This leads to situations where, in one of the rooms, the light is sufficiently close to the light so that a new causal relation ($light \rightarrow thermometer$)




Fig. 11: In the bathroom, the light is close enough to the thermometer to influence temperature measures. This relation is detected by our algorithm (below).

appears in the causal model of the room. Faced with this event, our algorithm was able to see the new connection in the corresponding room.

In this particular case, the diagnostic inference offered by our approach allows finding the cause of a particular behavior of the system (figure 11): “*The thermometer reports a hot temperature in the bathroom while the heater is off*”. In this case, diagnostic inference on the Bayesian network would initiate $T = 1$ and $H = 0$, and would infer the state of the lamp $P(L = 1 \mid (T, H) = (1, 0)) = 0.61$.

A final line of analysis would be to look at the evolution of the causal graph over time. Its dynamics could allow us to explain abnormal phenomena that will occur after it has changed.

VI. LIMITATIONS AND FUTURE WORK

While our experimental setup of the smart home was set to be close to a real-life use case, some limitations remain in our approach, some of which will be discussed here.

The first major assumption is that the system’s causal model can be represented by an acyclic graph. This limitation is common in the literature on Causality Theory ([21, 18, 25]). However, especially in the context of SAS, retro-actions may occur between devices and the variables they monitor: for instance, consider how a “smart” heater would turn on depending on the room’s temperature. One potential workaround would be to take time into account, which removes any ambiguity regarding the direction of a causal relation. This can be done by learning a *Temporal Bayesian Network* [5], which takes into account the time dependency.

However, adding time to the equation is no easy task: in sec. V, we argued that for our demonstration, we used a fixed time Δ_t after which we consider the consequences of interventions. This fixed value entails several issues: since different causal mechanisms can have different time characteristics, how can one know how long is long enough when waiting for the consequences of intervention? For future work, we may therefore consider implementing a multi-scale approach [6] to the environment model to cope with different time constants.

Furthermore, our approach requires a certain number of interventions on the system and has shown to perform better

when only a limited number of variables are non-doable. These issues are minor in a simulated example but can be limiting when operating in a real-life environment. A workaround is to consider having access to a model on the environment, such as a “digital twin” [23], which our algorithm can use.

Having the causal diagram of a system, as opposed to a simple Bayesian Network, offers possibilities to be integrated into explanation frameworks. For instance, we may use Causal Bayesian Network in conjunction with the general Explanatory Engine proposed by [7]. Diagnostic can be used as a powerful tool to propose hypotheses for abductive inference (i.e. finding the cause of an observed phenomenon), while predictive inference might be used to explore the potential consequences of a proposed solution, in the context of an explanation. The Causal Bayesian Network can also be converted into a Fuzzy Cognitive Map (FCM) which is used to present causal knowledge intuitively because of its simplicity. [28] propose a method to generate an FCM from the Bayesian network using Pearson’s correlation coefficient equation. Here we can use the scores of the causality test instead of the correlation.

Future work may also focus on learning for large networks, for example, use multi-scale learning. We mean by that learning the model of rooms and then going to a larger scale to learn the interactions between the rooms.

VII. CONCLUSION

This study is an example showing how some methods of Causality Theory can be implemented in the context of real-life SAS. We have seen how intervention operations can be simulated to train a causal diagram, which can then be used as a basis for our Causal Bayesian Network. This technique has shown encouraging results in the example of smart home. Since it required no ad hoc knowledge about the particularities of the smart home context, it can be generalized to other comparable settings.

Knowing the Causal Bayesian Graph offers advantages for applications such as explanations, given the more “natural” source of relations it entails, compared to a more traditional Bayesian Networks. We consider using this tool as a means to perform abductive and predictive inference in a broader explanatory framework for SAS.

REFERENCES

- [1] Elias Bareinboim et al. “On Pearl’s Hierarchy and the Foundations of Causal Inference”. In: *ACM special volume in honor of Judea Pearl* (2020).
- [2] Concha Bielza and Pedro Larranaga. “Discrete Bayesian network classifiers: A survey”. In: *ACM Computing Surveys (CSUR)* 47.1 (2014). Publisher: ACM New York, NY, USA.
- [3] Jianqing Fan, Fang Han, and Han Liu. “Challenges of big data analysis”. In: *National science review* 1.2 (2014). Publisher: Oxford University Press.
- [4] Koller & Friedman. “Probabilistic Graphical Models - Principles and Techniques”. In: 2009.

- [5] Pablo Hernandez-Leal et al. "Learning temporal nodes Bayesian networks". In: *International Journal of Approximate Reasoning* 54.8 (2013).
- [6] Mark F Horstemeyer. "Multiscale modeling: a review". In: *Practical aspects of computational chemistry*. Springer, 2009.
- [7] Etienne Houzé et al. "A Decentralized Approach to Explanatory Artificial Intelligence for Autonomic Systems". In: *ACSOS 2020 Conference Proceedings, Companion*. Aug. 2020.
- [8] Dominik Janzing and Bernhard Schölkopf. "Causal inference using the algorithmic Markov condition". In: *IEEE Transactions on Information Theory* 56.10 (2010). Publisher: IEEE.
- [9] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [10] Jeff Kramer and Jeff Magee. "A rigorous architectural approach to adaptive software engineering". In: *Journal of Computer Science and Technology* 24.2 (2009).
- [11] J Naveen Ananda Kumar and Srija Chimmani. "Proposal of smart home resource management for waste reduction and sustainability using AI and ML". In: *2019 International Conference on Communication and Electronics Systems (ICCES)*. IEEE, 2019.
- [12] Philippe Lalanda and Catherine Hamon. "A service-oriented edge platform for cyber-physical systems". In: *CCF Transactions on Pervasive Computing and Interaction* 2 (2020). Publisher: Springer.
- [13] Nianyu Li et al. "Explanations for human-on-the-loop: A probabilistic model checking approach". In: *Proceedings of the IEEE/ACM 15th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*. 2020.
- [14] Prashan Madumal et al. "Explainable reinforcement learning through a causal lens". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. Issue: 03. 2020.
- [15] Stijn Meganck, Philippe Leray, and Bernard Manderick. "Learning Causal Bayesian Networks from Observations and Experiments: A Decision Theoretic Approach". In: *Modeling Decisions for Artificial Intelligence*. Ed. by Vicenç Torra et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006.
- [16] Bojan Mihaljević, Concha Bielza, and Pedro Larrañaga. "Learning Bayesian network classifiers with completed partially directed acyclic graphs". In: *International Conference on Probabilistic Graphical Models*. PMLR, 2018.
- [17] Ziad Obermeyer et al. "Dissecting racial bias in an algorithm used to manage the health of populations". In: *Science* 366.6464 (2019).
- [18] Judea Pearl. *Causality: Models, Reasoning and Inference*. 2nd. USA: Cambridge University Press, 2009.
- [19] Judea Pearl. "Fusion, propagation, and structuring in belief networks". In: *Artificial intelligence* 29.3 (1986). Publisher: Elsevier.
- [20] Judea Pearl and Dana Mackenzie. *The Book of Why. The New Science of Cause and Effect*. New York: Basic Books, 2018.
- [21] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. MIT press, 2017.
- [22] Jonathan G Richens, Ciarán M Lee, and Saurabh Johri. "Improving the accuracy of medical diagnosis with causal machine learning". In: *Nature communications* 11.1 (2020). Publisher: Nature Publishing Group.
- [23] Roland Rosen et al. "About the importance of autonomy and digital twins for the future of manufacturing". In: *IFAC-PapersOnLine* 48.3 (2015). Publisher: Elsevier.
- [24] Julian Schrittwieser et al. "Mastering atari, go, chess and shogi by planning with a learned model". In: *Nature* 588.7839 (2020). Publisher: Nature Publishing Group.
- [25] Peter Spirtes et al. *Causation, prediction, and search*. MIT press, 2000.
- [26] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [27] Sergei Volodin, Nevan Wichers, and Jeremy Nixon. "Resolving spurious correlations in causal models of environments via interventions". In: *arXiv preprint arXiv:2002.05217* (2020).
- [28] Yit Yin Wee et al. "A method for root cause analysis with a Bayesian belief network and fuzzy cognitive map". In: *Expert Systems with Applications* 42.1 (2015).
- [29] Danny Weyns. "Software Engineering of Self-adaptive Systems". In: *Handbook of Software Engineering*. Ed. by Sungdeok Cha, Richard N. Taylor, and Kyochul Kang. Cham: Springer International Publishing, 2019.
- [30] Sewall Wright. "Correlation and causation". In: *Journal of agricultural Research* 20 (1921).
- [31] K Yang and Sung-Bae Cho. "A context-aware system in Internet of Things using modular Bayesian networks". In: *International Journal of Distributed Sensor Networks* 13 (May 2017).