Causal Inference in NARS

Bowen Xu, Pei Wang ✉

Department of Computer and Information Sciences,
Temple University, Philadelphia, USA
{bowen.xu, pei.wang}@temple.edu

Abstract. Humans engage in causal inference almost every day, however, the term 'causation' is still quite ambiguous, and few AI systems provide a comprehensive and satisfactory solution to causal inference. In this paper, we adopt the primary meaning of causation, *i.e.*, *prediction*, and argue that in different contexts other demands are attached to it. We describe the approach of causal inference in NARS and present some working examples, both at the sensorimotor and abstract levels. The theoretical and practical consequences are quite different from traditional AI approaches.

Keywords: Causal Inference · Non-Axiomatic Reasoning System · Prediction · Explanation

1 Introduction

Causal inference has attracted wide interest over the years in Artificial Intelligence (AI), though it has a rather long history in philosophy, psychology, economics, physics, and so on. On the one hand, causal inference seems to be a fundamental capability of intelligent agents to interact with the world; on the other hand, there is no consensus on what causation is, let alone the mechanism of causal inference [12]. Some people consider causation a fat concept with diverse meanings in different contexts [2, 3]. Nevertheless, over the years, some AI researchers (*e.g.* [10]) believed that causal inference would bring about a scientific revolution, attracting broad interest in enabling AI to make causal inferences.

Traditionally, causation is viewed as an objective relation between events, and various definitions were proposed from distinct perspectives (*e.g.*, regularity, counterfactualism, interventionism, *etc.*) [2,7,12]. In contrast, causation has been treated as a subjective relation in Non-Axiomatic Reasoning System (NARS) [6, 14] where no assumption is made on the "true cause(s)" of an event [18], though predicting the future or explaining the past based on the current situation is referred to as "causal inference".

In this paper, we further discuss the nature of causal inference in NARS. We start by briefly reviewing previous works, especially in AI, and then proposing our positions and perspectives on causation and causal inference. Finally, some working examples are provided to show that NARS can do causal inference in a unified manner: both in sensorimotor procedure and abstract-level reasoning.

2 Previous Works

Causation and causal inference have been discussed in many domains. For example, the representative of the subjective view of causation is David Hume, a philosopher who believes that it is impossible to deduce consistent laws from existing observations and that 'causation' is nothing but a human mental habit or even an illusion. In contrast, the representative of the objective view of causation is Pierre-Simon Laplace, who saw the present state of the universe as the result of its previous state and the cause of its subsequent state so that an omniscient and omnipotent intellect could accurately predict the future according to causal laws.

In Artificial Intelligence, the most influential approaches to causal inference are probabilistic, as exemplified by the *structural causal models* (SCM) proposed by Judea Pearl [9], where causes and effects are nodes in a graph, and directed connections represent causal relations. In Pearl's theory, a "causal model" is composed of a set of "causal assumptions" (or prior knowledge) provided by human scientists, and causal inference aims to evaluate the extent of the truthfulness of an implication. Pearl formalizes *intervention* using "*do*-calculus" [10], by which a conditional probability is adjusted to eliminate cognitive bias. However, an often-overlooked implicit assumption in Pearl's theory is that future observation will follow the same regularity as the past. With this assumption, given a causal model and collected data, an actual estimate of effects can be obtained as if it did some interventions to collect more data. As a result, one is able either to know "true effects" by mere observation or to know such estimation is unavailable. The same assumption is made in Pearl's theory on counterfactual.

Instead of assuming an available causal model, some researchers focus on discovering graph structures as possible causal explanations for data, and this ramification of research is called "causal discovery" [13, 19]. Solving a causal discovery problem was defined as "*recovering the true graph from the given data set*" [19], implicitly assuming the existence and availability of a *true* causal graph. In some cases, the aim of causal discovery is "*revealing causal information by analyzing purely observational data*" [4].

We argue that the traditional views on causal inference mentioned above do not apply to *AGI* under our conception [16]. It does not mean that their works are wrong, but that, as far as *causal inference* is concerned, AGI systems should work under different conditions from the previous ones. Firstly, AGI systems actively acquire data rather than passively observing, as opposed to causal inference [10] in the traditioal sense. Secondly, AGI systems learn causal relations among events by themselves without given knowledge. Thirdly, AGI systems may face situations where future observations are not guaranteed to conform to past findings, so that "true (causal) graphs" [19] should not be assumed in causal discovery. Finally, there cannot be a causal model containing all possible events, as many events are known to the system only after their occurrence.

While in many traditional models "causal inference" only covers *causal estimation*, in NARS that phrase also covers *causal discovery*, as the two functions

are carried out by the same mechanism, rather than handled separately as in the traditional models.

3 The NARS Approach

Because NARS is based on the *Assumption of Insufficient Knowledge and Resources*, the objective view of causation cannot be accepted [14]. There is no practical difference between saying "Objective causal laws do exist, but we can never get them" and saying "There are no objective laws". It is even impossible to say "Our knowledge is an approximation of objective laws" because such a statement makes sense only if the degree of approximation can be determined. The historical success of science does not mean that we can obtain concrete "objective causation", because no theory, however successful in the past, can logically guarantee the correctness of all its predictions for the future, as Hume argued [8]. But we are also against *agnosticism* because the result is often the abandonment of all effort and even the denial of the value of thinking. The key here is how to interpret the meaning of *knowing*. We believe the *knowledge* of a cognitive subject depends on the cognitive ability and mechanism of the subject, so it is improper to take "knowable" as meaning "be able to be obtained in a complete, accurate, and irrefutable description" of an object. All objects are knowable, but all knowledge is subject to revision or rejection.

In the case of causal knowledge, admitting that they are "mental habits" does not necessarily make them less valuable. The adaptability of intelligent systems requires effectively predicting the future and controlling it as much as possible. This requires finding regularities in experience, even though the results are not guaranteed to hold in the future, or they have known counter-evidence. The rationality in this situation is to be *as good as possible*, rather than *absolutely good*. "Causality" under this interpretation is not the objective connection between events (even what counts as an "event" and whether there is a "connection" depends on the cognitive subject), but the rational connection between events that the cognitive subject can use to make predictions. To try to find a "cause" for an event is to explore the possibility of predicting and controlling it. This is a basic requirement for adapting to the environment, and an intelligent system must strive to do so, no matter how well it can do it under the circumstances.

Based on this perspective, "causation" is not a *logical* relationship in NARS, but is usually represented as acquired implicational statements with temporal order and other domain-specific attributes. In the simplest case, the statement "A is the cause of B" can be represented as "$A \not\Rightarrow B$", where "$\Rightarrow$" is implication, and "$/$" indicates that A occurs before B. Similar to other types of knowledge in NARS, such causal knowledge can be acquired either directly from external input or derived through the system's reasoning activity.

One common way of deriving causation is "temporal induction", where the conclusion "A always leads to B when A happens before B" is generated from observed instances of "A happens, and then B happens". Like other inductive

conclusions, the initial *confidence* of such a conclusion is relatively low, which can be referred to as a "hypothesis". Only when similar observations are repeated numerous times, a reliable conclusion can be obtained. In contrast, the deductive relationship "$A \not\Rightarrow B$", derived from "$A \not\Rightarrow C$" and "$C \not\Rightarrow B$", has a higher *confidence* value.

The reason for causation to be represented in this form is that the original meaning of "$A \not\Rightarrow B$" is "prediction", that is, to infer "B is about to happen" from "A has happened", and this is the major role of causation in adaptation. Therefore, in NARS, there is no separate "causal reasoning" mechanism, but the temporal order of events is taken into account in the inference rules dealing with implication. This is because what we commonly call "causation" does not have a definite logical meaning as "$A \not\Rightarrow B$", but largely depends on contextual factors such as domain and purpose.

Based on this view, let's analyze the given options and determine which one is most suitable to be considered "the cause" of the collective increase in stock prices in a certain sector of a stock market –

1. increased demand for related products in the recent period,
2. technological breakthroughs suggesting the imminent release of new products,
3. the release of statistical data earlier that day enhanced investors' confidence, and
4. a large number of people buying stocks in that sector on that day.

Journalists may choose (3) in their coverage, while economists are more likely to choose (1) or (2) and start debating which factor is more important. They will all think that (4) is self-evident and therefore not worth mentioning. Since the logic behind these choices is "to infer that B is about to happen from the fact that A has already happened", it is suggested that "A is the cause of B". Besides the necessary factor of "prediction", there are several other relevant factors, although they are not taken into account in every context:

- **Necessity:** In some fields, A is a sufficient condition for B (*i.e.*, "$A \not\Rightarrow B$"), while in some other fields, A is a necessary and sufficient condition for B (*i.e.*, "$A \not\Leftrightarrow B$").
- **Reality:** In some fields, causation is meaningful only when both A and B have occurred, while in others, they can be hypothetical or imagined events (*e.g.*, "What caused me to get wet was that I did not have an umbrella").
- **Counterexamples:** In some fields, a single counterexample can reject a proposed causal relation, while in some other fields, a causal relationship can have exceptions.
- **Controllability:** Some contexts require that A must be an action of an agent or under its control to be considered the cause of B.
- **Reproducibility:** In certain domains, causal relations can only be established between repeatable events, rather than between one-time events.
- **Explanation:** In some fields, a correlation between A and B alone is not enough to establish a causal relation between them, and an explanation is required (*e.g.*, A causes M and M causes B).

- **Temporal Distance:** The temporal distance between A and B may affect whether A is considered the cause of B, otherwise, the "Big Bang" could have been the cause of everything.
- **Exclusivity:** In some contexts, there might be only one cause for an event, while in others, multiple causes may be accepted.

The requirements above are all reasonable, but none of them is universally required when "causation" is considered. It would be nice to meet all the requirements, but if there is no such knowledge that meets them all, "no prediction at all" could be worse than wrong predictions. Thus, NARS allows for the coexistence of different types of causation in various contexts, each of which can attach additional requirements besides "supporting prediction". As a result, the specific meaning of 'causation' needs to be learned in each domain (physics, medicine, geology, economics, law, *etc.*), which may have domain-specific conditions. For example, "Gravity is the cause of a ball falling" can be expressed in NARS as "$(gravity \times falling\text{-}ball) \rightarrow physical\text{-}cause$", where "*physical-cause*" is an acquired relation that represents "cause" in physics. In our opinion, trying to deal with causation with multiple interpretations in a single manner is the main cause of the confusion in current research.

According to Piaget's theory, the concept of causation in infants begins to form through their understanding of the consequences of their actions [11]. In NARS, it can be represented as "$(S_1, \Uparrow Op) \Rightarrow S_2$", which means "After seeing S_1, if I perform operation $\Uparrow Op$, I will see S_2". Through analogy and inductive reasoning, it can be generalized to "$(S_1, S_3) \Rightarrow S_2$", which means "After seeing S_1, if any person performs the operation S_3 (which corresponds to my $\Uparrow Op$), the person will see S_2", where "the operation corresponding to my $\Uparrow Op$" is considered as an event S_3 (since it is not performed by me). This relationship can be further generalized to "$S_4 \Rightarrow S_2$", where S_4 is the name of the compound event '(S_1, S_3). In this form, there is no need for any cognitive subject to act as the cause of a result. Although a mature cognitive system may acquire causal knowledge directly from communicating with other systems, without going through the process of "from self to other agents then to events", this explanation can still apply to many phenomena related to causal inference. Note that the procedure of causal inference above also involves self-awareness [1], but due to the word limit, self-awareness is not discussed in this paper (refer to [17] for further information).

A direct consequence of this view on causation is that no intelligent system can accurately foresee all the consequences of an event, including its actions. As the saying goes, "What seems like a loss may be a gain." However, this is not a reason to refrain from predicting or taking action, and all decisions should be based on weighing the pros and cons of foreseeable consequences.

4 Working Examples

In the following, several working examples of causal inference in NARS are described and analyzed.

4.1 Sensorimotor integration

An example is related to the sensorimotor procedure, aiming to play the CartPole game in reinforcement learning. An agent should learn some general rules, such as that if the pole leans to the right, the cart will move to the right. An implementation of NARS [5] is used to play the game. The performance of the system is shown in Fig. 1, where the score indicates the proportion of time the system successfully balanced the pole over the past 100 moments. By checking its memory, the following beliefs used in decision-making are acquired by the system

$$((angle1, \Uparrow left) \Rightarrow good).$$
$$((angle2, \Uparrow right) \Rightarrow good).$$

These are causal beliefs about the environment, stating under what condition, by taking what operation, what effect would be caused. Based on the causal beliefs, NARS learns to perform satisfactorily in this environment.

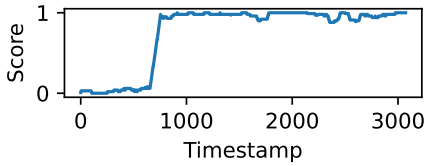


Fig. 1. The performance of NARS playing the CartPole game.

4.2 Simpson's paradox

Besides the sensorimotor procedure, NARS can do causal inference in a more abstract level, following the same logic. An example of intervention analyzed by [10] is Simpson's paradox. However, in NARS, quite different conclusions are drawn from Pearl's theory of causal inference. For example (see Tab. 1), the rates of heart attacks with drugs are 7.5% for women and 40% for men, and the rates without drugs are 5.3% for women and 30% for men. However, the rates of heart attacks for all are 18.3% with the drug and 21.7% without the drug, that is, the drug increases the risk of heart attack for men, and it also increases the risk for women, but it decreases the risk for the population as a whole!

	Control Group (No Drug)		Treatment Group (Took Drug)	
	Heart Attack	No Heart Attack	Heart Attack	No Heart Attack
Female	1	19	3	37
Male	12	28	8	12
Total	13	47	11	49

Table 1. Fictitious data illustrating Simpson's paradox. This table stems from [10] (Table 6.4).

In [10], Pearl explained that this is mainly because gender is a "confounder", and a person with different genders may have different selections of whether to

take the drug, and the risk of heart attack is also directly influenced by gender. They suggested to do statistical adjustment by using the back-door criterion, as a result, $P(H|do(D)) = 0.238$, and $P(H|do(\neg D)) = 0.177$. In this way, they said the correct conclusion can be drawn that the drug still harms the whole population. Despite the number game behind the paradox, let us check the implicit assumptions behind this statistical adjustment. The interpretation of *do*-calculus in Pearl's causal inference theory is that a certain intervention is executed without really doing it so that one can evaluate the "true causal effect" on observed data without further experiments. It is assumed that the data are perfect and sufficient so that regularity would hold if more observations were acquired by action.

In NARS, we have a different explanation and argue that the two conflict beliefs are reasonable to occur in the human mind. Given the two premises

$$((male \wedge D) \nRightarrow H).\langle 0.4; 0.95\rangle$$
$$((female \wedge D) \nRightarrow H).\langle 0.08; 0.98\rangle$$

and the compositional rule[1], it is derived that

$$(((male \vee female) \wedge D) \nRightarrow H).\langle 0.03; 0.93\rangle$$

Similarly, we have "$(((male \vee female) \wedge \neg D) \nRightarrow H).\langle 0.02; 0.93\rangle$". It shows that the frequency of people who took the drug and then suffered heart attacks is smaller than that of people who did not take the drug, *i.e.*, the drug increases the risk of heart attacks. In contrast, from direct observations,

$$(D \nRightarrow H).\langle 0.18; 0.98\rangle$$
$$(\neg D \nRightarrow H).\langle 0.28; 0.98\rangle$$

taking the drug seems better. The two contradictory beliefs coexist in NARS' mind, thus it is a paradox from its view.

4.3 Intervention & Counterfactual

To better illustrate causal inference in NARS, the example shown in Fig. 2 can be used. Suppose that the relevant events mentioned in the following have been obtained from perception: a remote control, C, can be manipulated by a system (like NARS, or a human being) who wants to watch TV shows. Once the system presses a button on C, the corresponding event (as an operation) $\Uparrow P$ occurs in its mind. Subsequently, the indicator light flashes once (event L), and then the TV is turned on and displays a picture (event D).

The system constructs the knowledge "$(\Uparrow P \nRightarrow D)$" through *temporal induction*: given events $\Uparrow P$ and D as well as their occurrence time, "$(\Uparrow P \nRightarrow D)$" is implied. The *confidence* value is low at first (as it has been tested only once), so

[1] $\{T_1 \Rightarrow M, T_2 \Rightarrow M\} \vdash (T_1 \vee T_2) \Rightarrow M\langle F_{int}\rangle$ [15]. It is isomorphic in temporal inference.

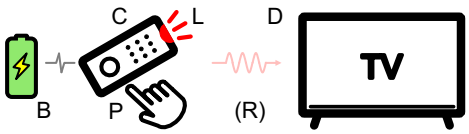


Fig. 2. An Example of Causal Inference

it can be viewed as a hypothesis, and the *confidence* value gets higher as more evidence is observed. In most cases, this knowledge is more than enough for the system to predict the future and explain the past, though similarly, it can also form an "explanation chain" for "$\Uparrow P$ causes D": it holds because "$\Uparrow P$ causes L" and "L causes D", that is, "$(\Uparrow P \not\Rightarrow L)$" and "$((\Uparrow P, L) \not\Rightarrow D)$".

The system usually considers $\Uparrow P$ as the cause of L, rather than "$(B, \Uparrow P)$" (where B indicates the existence of a battery in C) because the former is simpler than the latter and the battery is usually there; however, if $\Uparrow P$ by is not followed by L as the battery is missing, "$(B, \Uparrow P)$" is better to be viewed as the cause, because it is more correct. Similarly, when $\Uparrow P$ occurs but D does not subsequently occur, "$(\Uparrow P, L)$" is better to be explained as the cause of D.

There is no "true cause" of TV's displaying D, but there are satisfactory explanations for it. Infants may never generate explanations "Pressing the button $\Uparrow P$ causes an infrared ray R" and "The infrared ray R causes displaying D"; similarly, ill-educated adults may not know "$\Uparrow P$ causes electric current", "the current causes electronic transitions in an infrared emitting diodes", and "electronic transitions produce infrared ray R"... There can be an extremely long chain for explanation. Nevertheless, in daily life, satisfactory explanations are enough to be viewed as causal relations.

Intervention: *Intervention* is nothing but *operation* in NARS. In the context of causal inference, *operation* as *intervention* plays two roles – actively acquiring new knowledge, and filtering competing hypotheses in decision-making. By *operations*, the system changes the environment to construct new knowledge or to collect more evidence for existing knowledge, and some knowledge might not be easy to obtain by pure observation. In this example, the system cannot know "$((\Uparrow P, L) \not\Rightarrow D)$" until actually executing the operation $\Uparrow P$.

Counterfactual: Through exploration, the system finds that "battery" is necessary for L, *i.e.*, "$(B, \Uparrow P) \not\Rightarrow L$"; the battery can be installed through the action "$\Uparrow install$", *i.e.*, "$(\Uparrow install \not\Rightarrow B)$". However, when the system presses the button (*i.e.*, $\Uparrow P$) and the light does not work as anticipated (*i.e.*, $\neg L$), the system might observe the lack of battery (*i.e.*, $\neg B$). Since "if B were true, as a *counterfactual*, then L would also be true", and installing the battery (*i.e.*, $\Uparrow install$) implies B, it is derived that if the system installed the battery, the light would flash, that is, "$((\Uparrow install, \Uparrow P) \not\Rightarrow L)$". This is a preliminary form of *counterfactual* reasoning in NARS, although the related inference rules have existed in the logic [15].

Counterfactual reasoning exactly means reasoning with a statement that is contrary to the fact. In this example, $\neg B$ is the *fact* (according to its *truth value*), while "$(B, \Uparrow P) \not\Rightarrow L$" is a basis for concluding of "what if B".

More generally, suppose the system has a belief "$(\Uparrow assume(B) \not\Rightarrow B)$", and the system presses the button, then an implication is "$(\Uparrow assume(B) \not\Rightarrow L)$", and other conclusions can be drawn with the precondition "$\Uparrow assume(B)$", such as "$(\Uparrow assume(B) \not\Rightarrow D)$". It means that, even if event D does not actually occur, the system knows "it would have occurred if the remote had the battery".

5 Conclusion

In this paper, we focus on the issue of causal inference in NARS. We argue that the concept "causation" has different meanings in different contexts or disciplines, while the intersection among them is "prediction". Additional demands are attached to "prediction" in various contexts.

In NARS, a *cause* is a kind of *explanation* (namely, *causal explanation*) of an *event*. *Causal explanation* is one of various types of *explanations*, while a *cause* is selected among multiple premises of an *event* (as the conclusion) according to different contexts. The *event* as the conclusion here is also named *effect*. *Causal relation* is the relation between a *cause* and a corresponding *effect*, however, it is not a built-in logical relation, since it has quite different meanings under different contexts and domains. Nevertheless, the original form of causal relation is prediction, and in a specific context, it is a type of *acquired relation*.

We present some working examples of causal inference in NARS, though they do not involve inferring with *acquired relations*. We show that causal inference happens at both the sensorimotor level and the abstract level: the system learns some predictive implications, which can be viewed as *conditioned reflex*, from sensory signals; we address the issue of Simpson's paradox and argue that the cognitive bias also occurs in NARS, and this is reasonable under the *Assumption of Insufficient Knowledge and Resources*; we design a circumstance showing intervention and counterfactual reasoning happen in NARS. By intervention, the system actively obtained new evidence from the environment, so that new knowledge is acquired, which cannot be done by pure observation. Intervention also helps to filter competing hypotheses in decision-making, but we do not discuss this aspect in detail for lack of space. Counterfactual reasoning is about reasoning on the statements that are *false* actually, deriving possible consequences as if those were *true*.

The philosophical positions on causation and the practice of causal inference in NARS are quite different from the previous works. This work is preliminary, nonetheless, we suggest that it is a potential and promising solution to causal inference in AGI.

Acknowledgement. We sincerely thank Patrick Hammer's implementation of NARS, which helped for the experimental part of this paper. We also appreciate the anonymous reviewers for their valuable feedback.

References

1. Bennett, M.T.: Emergent Causality and the Foundation of Consciousness. In: Hammer, P., Alirezaie, M., Strannegård, C. (eds.) Artificial General Intelligence, vol. 13921, pp. 52–61. Springer Nature Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-33469-6$_6$
2. Broadbent, A.: Causation, `https://iep.utm.edu/causation/`, internet Encyclopedia of Philosophy. URL: https://iep.utm.edu/causation/. Accessed date: 2023-06-26
3. Cartwright, N.: Causation: One Word, Many Things. Philosophy of Science **71**(5), 805–819 (Dec 2004). https://doi.org/10.1086/426771, `https://www.cambridge.org/core/product/identifier/S0031824800002889/type/journal_article`
4. Glymour, C., Zhang, K., Spirtes, P.: Review of Causal Discovery Methods Based on Graphical Models. Frontiers in Genetics **10** (2019), `https://www.frontiersin.org/articles/10.3389/fgene.2019.00524`
5. Hammer, P., Lofthouse, T., Fenoglio, E., Latapie, H., Wang, P.: A reasoning based model for anomaly detection in the smart city domain. In: Arai, K., Kapoor, S., Bhatia, R. (eds.) Intelligent Systems and Applications. pp. 144–159. Springer International Publishing, Cham (2021)
6. Hammer, P., Lofthouse, T., Wang, P.: The OpenNARS implementation of the Non-Axiomatic Reasoning System. In: Steunebrink, B., Wang, P., Goertzel, B. (eds.) Proceedings of the Ninth Conference on Artificial General Intelligence. pp. 160–170 (2016)
7. Hitchcock, C.: Causal Models. In: Zalta, E.N., Nodelman, U. (eds.) The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, spring 2023 edn. (2023), `https://plato.stanford.edu/archives/spr2023/entries/causal-models/`
8. Hume, D.: An Enquiry Concerning Human Understanding. London (1748)
9. Pearl, J.: Causal inference in statistics: An overview. Statistics Surveys **3**, 96–146 (Jan 2009). https://doi.org/10.1214/09-SS057
10. Pearl, J., Mackenzie, D.: The Book of Why. Basic Books, New York (2018)
11. Piaget, J.: The construction of reality in the child. Basic Books, New York (1954)
12. Schrenk, M.: Metaphysics of Science: A Systematic and Historical Introduction. Routledge (Aug 2016)
13. Spirtes, P., Glymour, C.N., Scheines, R.: Causation, Prediction, and Search. MIT Press (2000)
14. Wang, P.: Non-Axiomatic Reasoning System: Exploring the Essence of Intelligence. Ph.D. thesis, Indiana University (1995)
15. Wang, P.: Non-Axiomatic Logic: A Model of Intelligent Reasoning. World Scientific, Singapore (2013)
16. Wang, P.: On defining artificial intelligence. Journal of Artificial General Intelligence **10**(2), 1–37 (2019). https://doi.org/10.2478/jagi-2019-0002
17. Wang, P.: A constructive explanation of consciousness. Journal of Artificial Intelligence and Consciousness **7**(2), 257–275 (2020)
18. Wang, P., Hammer, P.: Issues in temporal and causal inference. In: Bieger, J., Goertzel, B., Potapov, A. (eds.) Proceedings of the Eighth Conference on Artificial General Intelligence. pp. 208–217 (2015)
19. Zanga, A., Ozkirimli, E., Stella, F.: A survey on causal discovery: Theory and practice. International Journal of Approximate Reasoning **151**, 101–129 (2022)