



**HAL**  
open science

# Action Languages Based Actual Causality in Ethical Decision Making Contexts

Camilo Sarmiento, Gauvain Bourgne, Daniele Cavalli, Katsumi Inoue, Jean-Gabriel Ganascia

## ► To cite this version:

Camilo Sarmiento, Gauvain Bourgne, Daniele Cavalli, Katsumi Inoue, Jean-Gabriel Ganascia. Action Languages Based Actual Causality in Ethical Decision Making Contexts. 2022. <hal-03790699>

**HAL Id: hal-03790699**

**<https://hal.science/hal-03790699v1>**

Preprint submitted on 28 Sep 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-SA 4.0 - Attribution - ShareAlike - International License

---

# ACTION LANGUAGES BASED ACTUAL CAUSALITY IN ETHICAL DECISION MAKING CONTEXTS

---

**Camilo Sarmiento**

Sorbonne Université, CNRS, LIP6  
F-75005 Paris, France  
camilo.sarmiento@lip6.fr

**Gauvain Bourgne**

Sorbonne Université, CNRS, LIP6  
F-75005 Paris, France  
gauvain.bourgne@lip6.fr

**Daniele Cavalli**

École Normale Supérieure -  
PSL Research University,  
F-75005 Paris, France  
daniele.cavalli@ens.psl.eu

**Katsumi Inoue**

National Institute of Informatics  
Tokyo, Japan  
inoue@nii.ac.jp

**Jean-Gabriel Ganascia**

Sorbonne Université, CNRS, LIP6  
F-75005 Paris, France  
jean-gabriel.ganascia@lip6.fr

## ABSTRACT

Moral responsibility is closely intermixed with causality, even if it cannot be reduced to it. Besides, rationally understanding the evolution of the physical world is inherently linked with the idea of causality. It follows that decision making applications based on automated planning, especially if they integrate references to ethical norms, have inevitably to deal with causality. Despite these considerations, much of the work in computational ethics relegates causality to the background, if not ignores it completely. This paper contribution is double. The first one is to link up two research topics—automated planning and causality—by proposing an actual causation definition suitable for action languages. This definition is a formalisation of Wright’s NESS test of causation. The second is to link up computational ethics and causality by showing the importance of causality in the simulation of ethical reasoning and by enabling the domain to deal with situations that were previously out of reach thanks to the actual causation definition proposed.

**Keywords** Computational Ethics · Causality · Actual Causality · Regularity Theories of Causation · Action Languages

## 1 Introduction

We aim to design an ethical supervisor that can be embedded in agents so that their actions obey several moral prescriptions. The risks associated with their actions are then limited. The need for such a supervisor arises because of the growing number of agents capable of handling more complex tasks to whom more and more responsibilities are delegated [Tolmeijer et al., 2021]. As a result, their place in our daily lives is growing, as is the risk associated with this trend. The agents we are interested in are

technical objects capable of: taking into account the information they receive from the environment and acting to modify it. This definition should not be confused with the one used in the philosophy of action [Schlosser, 2019] and human science in general, where the concept of agent is usually linked to agency. The ethical supervisor that we use is the  $\text{ACE}$  (Action-Causality-Ethics) modular framework presented in Figure 1. It is the logical continuation of a series of works [Berreby et al., 2015, 2017, 2018, Bourgne et al., 2021].




Figure 1: Modular framework for representing and applying ethical principles.

The purpose of this paper is first to defend the idea that if we want to design such a supervisor, it is essential to explore imputability—‘no blame or praise may be assigned without some account of causal relationship between an agent and an outcome’ [Beebe et al., 2009]. Logically, the second purpose of this paper is to propose a definition of causality that can be used to make decisions. This inevitably leads us to look at causation which will be the main focus of this paper. Because of its essential role in human reasoning—both in trivial and in complex situations—numerous works in a variety of disciplines have unsuccessfully tried to propose a widely agreed upon theory of causation. Since we are in an operational framework given that our focus is on ethical decision making, we can make a couple of assumptions while remaining relevant. Therefore, we place ourselves in a classical planning framework which assumes problems are discrete and deterministic. Unlike *type causality* which seeks to determine general causal relationships, *actual causality* fits our purpose because it is concerned with particular events [Halpern, 2016]. Limiting ourselves to a simplified framework and to actual causality does not make causality trivial, many issues remain.

This paper is structured as follows. Section 2 shows how important the ability to establish causal relationships is when modelling ethical reasoning. Despite encouraging initial results in the field of computational ethics, much remains to be done. We argue that this distance is partly due to limitations that can be explained by the lack of work incorporating mechanisms for establishing complex causal relationships. Section 3 introduces the action language semantics—in which we encode causal knowledge—allowing concurrency of events. Section 4 discusses what is the appropriate approach to causation for our causal inquiry. In this section we explore two main highly influential theories of causation: *regularity* and *counterfactual*. Section 5 offers a description of our actual causality definition proposal for our ethical supervisor. This definition is factual and independent of policy choices and allows to handle complex cases of causality in ethical decision making applications. Finally, we conclude and give some perspectives in Section 6.

## 2 Causality in Computational Ethics

In recent years, various works in the field of computational ethics have shown that it is possible to formalise moral prescriptions—and this by means of a large number of techniques [Tolmeijer et al., 2021]. Primarily, those belonging to the consequentialist and deontological traditions, which allow a better formalization than more recent theoretical reflections (e.g ethics of care, quantum ethics). To do this, these works have mainly relied on thought experiments, such as the trolley problem introduced by Foot [1967]. The first challenge was to formalise the problem—to represent it in a computer language so as to reproduce the sequence of the thought experiment. For the trolley problem, it was necessary to represent the fact that inaction would cause the death of five people, while the action of diverting the trolley would only cause the death of one person.

The second challenge was to formalise moral prescriptions for assessing the acceptability of possible action choices in the thought experiment. Various moral prescriptions have been formalised so far. Next are some of these prescriptions belonging to consequentialism with the works where they have been formalised: *prohibiting purely detrimental actions*, an action is impermissible if its consequences are purely bad [Berreby et al., 2017]; *principle of benefits vs. costs*, an action is impermissible if its bad consequences are more significant than its good consequences [Berreby et al., 2017, Bourgne et al., 2021]; *principle of least bad consequence*, an action is impermissible if its worst consequence is considered the least good of all possible alternative consequences [Ganascia, 2015, Berreby et al., 2017, Lindner et al., 2017]; *Pareto principle*, an action is impermissible if there is another possible action whose positive consequences are better or whose negative consequences are less bad [Lindner et al., 2017]; *act utilitarianism*, an action is impermissible if there is another action with consequences that are better overall [Berreby et al., 2017, Lindner et al., 2017, Bourgne et al., 2021]; *rule utilitarianism*, an action is impermissible if there is another action for which the universalization of the rule to which it conforms produces more utility than the universalization of the first action corresponding rule [Berreby et al., 2017]. Following are some of the moral prescriptions belonging to deontology with the works where they have been formalised: *codes of conduct*, an action is impermissible if it goes against a prohibition in the code followed [Berreby et al., 2017, Limarga et al., 2020]; *Kant's formula of the end in itself* [Kant, 2013], an action is impermissible if it has consequences—which involve an individual—other than the ends of the action [Berreby et al., 2017, Lindner and Bentzen, 2018, Bourgne et al., 2021]; *d'Aquin's doctrine of double effect* [d'Aquin, 1266], an action is impermissible if at least one of these four conditions is met: (i) the action is intrinsically bad, (ii) at least one bad consequence is a means to a good one, (iii) at least one bad consequence was intended, (iv) its bad consequences are greater than its good consequences [Berreby et al., 2015, 2017, Govindarajulu and Bringsjord, 2017, Lindner et al., 2017, Bonnemains et al., 2018]. The diversity represented by these few examples is proof that the formalisation of moral prescriptions is possible and not limited to a small family of ethical systems.

Despite these encouraging results, there is still much room for improvement. Indeed, what makes ethical reasoning so laborious to reproduce is that ethics lies in the subtlety of each problem. Thus, the complexity of the task facing computational ethics is as much in the formalisation of moral prescriptions as it is in the formalisation of the problem—more specifically in the integration of all the subtleties relevant to the problem into the formalisation that is made of it. Consider the situation where an agent has to decide whether the action of making a vaccine compulsory is permissible. Seen in this way, the conflict to be resolved will be primarily between the right to personal liberty and the right to collective safety [Assembly, 1948]. Undoubtedly, the ethical reasoning will change if a significant part of the population is allergic to the vaccine. Also, the ethical considerations will change radically if the group of allergic people is mainly composed of individuals of the same gender or ethnicity. If any agent is to make this decision based on an ethical decision making mechanism, it must be able to take into account the subtleties relevant to each of these cases. The solutions proposed at the moment are still far from being able to handle the complexity involved in the real world.

The main argument in this section is that the current limitations of existing solutions are mainly due to the systematic absence—with a few exceptions [Berreby et al., 2015, 2017, Bourgne et al., 2021, Lindner et al., 2017, Lindner and Bentzen, 2018]—of a mechanism for establishing causal relationships. The importance of causality appears obvious when the moral prescriptions belong to consequentialism, i.e. the theory that the acceptability of actions depends exclusively on the value of its consequences. However obvious, we shall see that causality is rarely given the place it deserves. Although the link seems less intuitive, causality is also essential in some theories of deontological ethics. This is the case when considering whether an action is a means to an end in Kant's second formulation of the categorical imperative or d'Aquin's doctrine of the double effect. Because of its importance in most of the formalised moral prescriptions, the absence of causality has two main consequences on the proposals made so far: (i) the oversimplification of the problem formalisation and (ii) the fact of leaving out problems—as Examples 1, 2, and 4—that may contain overdetermination [Wright, 1985]:

*overdetermined causation*: cases in which a factor other than the specified act would have been sufficient to produce the injury in the absence of the specified act, but its effects either (1) were preempted by the more immediately operative effects of the specified act or (2) combined with or duplicated those of the specified act to jointly produce the injury.

(i) The oversimplification of the problem arises from the lack of ability to create complex causal links. Hence, the only way to link an action to a consequence is either to formalise it as an intrinsic effect of the action, or to assess an initial and an end state and attribute all changes to the action. Considering the trolley problem, the first approach corresponds to formalising the action ‘pulling the switch lever’ as having the intrinsic effect of killing the individual on the second track while the action ‘doing nothing’ results in the death of the five individuals on the main track [Limarga et al., 2020, Bonnemains et al., 2018]. The second approach takes the initial state where all individuals are alive, compares it to the final state where some individuals are dead, and considers that all changes are effects of one of the two possible actions performed in between. These shortcuts have appeared to be viable solutions because the problems studied are mostly thought experiments—designed to help us understand the asymmetry of our judgement in certain situations, such as why it would be acceptable to sacrifice one person to save five in one context but not in another that is very similar [Nyholm and Smids, 2016]. Accordingly, in the studied problems the number of factors that can be taken into consideration is very small, the outcomes are certain, and the number of possible actions is very limited, a context far from that of the real world. This apparent simplicity has therefore masked the complexity of one of the challenges of the field: capturing the subtlety of each problem in its formalisation. These shortcuts are particularly inappropriate as they may have a negative influence on the ethical assessment of the problem. Given the first approach formalisation of the two possible actions in the trolley problem, these could be considered intrinsically bad in relation to the right to live because one of their intrinsic effects is the death of one or more individuals. The act of ‘pulling the switch lever’ would be in some way equated to shooting someone. The reader will quickly understand that this is not really the case and that it is a product of the formalisation. The second approach does not have this problem. However, it can be questioned about its capacity to manage imputability properly. Consider a case were one of the individuals in the trolley problem is shot by a hitman before the train diverts. The attribution of this death to the agent being able to divert the trolley is problematic. Yet, this is what happens if we reduce causality to a simple comparison of states. These imputation errors become unavoidable when other agents can act on the world and we consider possible to have concurrent actions—situations that are far from being exceptional and that we cannot simply ignore. The solution to both approaches undesirable effects is to formalise the dynamics of the problem. For instance, it would be more appropriate to formalise the two actions of the trolley problem as having an intrinsic effect on the movement of the trolley and not on the lives of the individuals. However, the agent’s actions are no longer directly linked to the death of the individuals in this setup. Thus, a mechanism for establishing causal relationships becomes essential.

(ii) The reason for not dealing with overdetermination is the difficulties that arise from it and that can only be managed by examining causality. Cases of overdetermination are commonly found in the complexity of reality (cause of pollution, cause of suicide, cause of economic loss, . . .), they give rise to numerous questions in law. They have therefore been the subject of much work in the fields of causation, whether by lawyers, philosophers or mathematicians. Insofar as overdetermination is a phenomenon that an agent in a real context may face, the inability of all the propositions mentioned so far to deal with these cases is an important limitation of the field. In addition, to avoid these cases, propositions leave aside concurrency of events which are also common cases an agent may face. Thus, to achieve the desired goal it is not only necessary that the proposals in the field incorporate a mechanism for establishing causal relationships, but also that these mechanisms are sufficiently complex to handle cases of causal overdetermination.

The examples below will be used throughout this paper. They present different cases of causal overdetermination while illustrating ethical issues that may arise and that we may wish to explore.

**Example 1** (*pollution—preemptive causation*). *A village along a river is home to  $n$  families. The drinking water used by the inhabitants of the village comes from a water treatment plant that draws water from the river, which in turn comes from a lake located upstream of the river. However, the capacity of this plant is limited, it can only treat water if it has a pollution indicator below a threshold. When water from the mountain reaches the lake, the pollution indicator is zero. There are two potential sources of pollution to the lake: (i) industrial wastewater from a factory that produces connected speakers for a famous online shopping site ( $w_s$ ) and (ii) industrial wastewater from a factory that produces life-saving medicines for  $k$  patients ( $w_m$ ). Under normal circumstances,  $w_m$  does not pollute the lake because it is treated by a wastewater treatment plant before being discharged into the lake. We assume that the discharges are necessary for the launching operation of the factory and that the production management from both factories is ensured by automated agents. We consider the scenario where the agent in charge*

*of the connected speakers factory launches the production. The discharged wastewater increases the pollution indicator to the threshold. It turns out that the plant treating the wastewater from the medicine factory has been out of order since the beginning of the month. A few hours after the launching of the connected speakers factory, the agent in charge of the medicines factory launches the production. The discharged wastewater increases the pollution indicator to two times the threshold. The inhabitants of the village are left without water in this scenario.*

Example 1 is a case of preemptive causation. The production of speakers and of medicines—given the broken state of  $w_m$  treatment plant—are both sufficient to cause the harm. Indeed, individually each discharge would raise the level to the threshold at which the village inhabitants are left without water. However, the eventuality that  $w_m$  discharges deprive the inhabitants of water is preempted by the anteriority of  $w_s$  discharges. In other words, effects of  $w_s$  took precedence over those of  $w_m$ . It is therefore the production of speakers that is identified as a cause of the harm.

**Example 2** (*pollution—duplicative causation*). *We remain in the same framework as the example 1, except that in this scenario the two agents launch the production at the same time. The discharged wastewater increases the pollution indicator to two times the threshold. The inhabitants of the village along the river are also left without water in this scenario.*

Example 2 is a case of duplicative causation. We remain in the same situation as above, except that this time the effects of  $w_s$  and  $w_m$  discharges on water quality occur simultaneously—they combine to produce the harm jointly. The production of speakers and medicines are both identified as causes of the harm.

Besides illustrating two types of overdetermination, these examples may well show the complexity behind ethical assessment. Indeed, many factors make the problem more complex. First of all, we have the fact that the value produced by the polluting activity is different. On the one hand the production of a so-called comfort good, on the other the production of a medicine essential for the survival of individuals. We could add that the speakers factory supports the region by giving work to a large part of the inhabitants of the village, whereas the medicine factory has completely automated its production line. We could also add that the discharges from one factory are authorised, while those from the other are not. In the end, all these factors can have a significant influence on the permissibility of each action. However, all this richness can only be explored if one has sufficient expressiveness to formalise it and a mechanism to establish sufficiently complex causal relationships. We will start by giving our proposal to address the first.

### 3 Action Language Semantics

The whole purpose of an action language is to determine the evolution of the world given a set of actions corresponding to deliberate choices of the agent. Those actions might trigger some chain reaction through external events. As a result, we need to keep track of both: the state of the world and the occurrence of events—the term ‘event’ connoting ‘the possibility of agentless actions’ [Russell and Norvig, 2010, chap 12]. The advanced state of maturity of PDDL [Ghallab et al., 1998, Haslum et al., 2019], its vocation to facilitate interchangeability, and its use by a large community, are all meaningful arguments in favour of this formalism—gradually extended by different fragments. However, the semantics of its deterministic fragment—corresponding to ADL [Pednault, 1989]—does not allow concurrency of events. To have a semantics that takes into account concurrency it is necessary to jump directly to PDDL+ [Fox and Long, 2006] which semantics is adapted to durative actions, thus inconsistent with our discrete time assumption. We therefore base our approach on an action language whose semantics is an intermediate point between the deterministic fragment of PDDL and PDDL+. This formalism works on a decomposition of the world into two sets:  $\mathbb{F}$  corresponding to variables describing the state of the world, more precisely ground fluents representing time-varying properties;  $\mathbb{E}$  representing variables describing transitions, more precisely ground events that modify fluents.

A fluent literal is either a fluent  $f \in \mathbb{F}$ , or its negation  $\neg f$ . We denote by  $Lit_{\mathbb{F}}$  the set of fluent literals in  $\mathbb{F}$ , where  $Lit_{\mathbb{F}} = \mathbb{F} \cup \{\neg f | f \in \mathbb{F}\}$ . The complement of a fluent literal  $l$  is defined as  $\bar{l} = \neg f$  if  $l = f$  or  $\bar{l} = f$  if  $l = \neg f$ . By extension, for a set  $L \subseteq Lit_{\mathbb{F}}$ , we have  $\bar{L} = \{\bar{l}, l \in L\}$ .

**Definition 1** (state  $s$ ). *The set  $L \subseteq Lit_{\mathbb{F}}$  is a state iff it is:*

- *Coherent*:  $\forall l \in L, \bar{l} \notin L$ .
- *Complete*:  $|L| = |\mathbb{F}|$ , i.e.  $\forall f \in \mathbb{F}, f \in L \text{ or } \neg f \in L$ .

A complete and coherent set of fluent literals thus determines the value of each of the fluents. An incoherent set cannot describe a reality. However, in the absence of information or for the sake of simplification, we can describe a problem through a coherent but incomplete set. We will call such a set a partial state. We model time linearly and in a discretised way to associate a state  $S_t$  to each time point  $t$  of a set  $\mathbb{T} = \{-1, 0, \dots, N\}$ . Having a bounded past formalisation of a real problem, we gather all states before  $t = 0$ —time point to which corresponds the state  $S_0$  that we call initial state—in an empty state  $S_{-1} = \emptyset$ .

We place ourselves in a framework of concurrency where  $E_t$  is the set of all events which occur at a time point  $t$ . Therefore,  $E_t$  is what generates the transition between the states  $S_t$  and  $S_{t+1}$ . Thus, the states follow one another as events occur, simulating the evolution of the world.  $E_{-1}$  is the set that gathers all events which took place before  $t = 0$ , such that  $E_{-1} = \{ini_l, l \in S_0\}$ . Events are characterised by two elements: preconditions give the conditions that must be satisfied by the state in order for them to take place; effects indicate the changes to the fluents that are expected to happen if they occur. The preconditions and effects are respectively represented as formulas of the language  $\mathcal{P}$  and  $\mathcal{E}$  defined as follows:

$$\mathcal{P} ::= l | \psi_1 \wedge \psi_2 | \psi_1 \vee \psi_2 \quad \mathcal{E} ::= [\psi]l | \varphi_1 \wedge \varphi_2$$

where  $l \in Lit_{\mathbb{F}}$ ,  $[\psi]l$  is the notation for the conditional effect indicating that  $l$  is an effect if the condition  $\psi$  is satisfied— $[\top]l$  is just written  $l$ —and the logical connectives  $\wedge, \vee$  have standard first-order semantics. We can then deduce that if  $\varphi \in \mathcal{E}$ ,  $\varphi = \bigwedge_{i \in 1, \dots, m} [\psi_i]l_i$ . For the sake of brevity, we adopt a set notation for  $\varphi \in \mathcal{E}$  which we will use where relevant, such that  $\varphi = \{[\psi_i]l_i, i \in 1, \dots, m\}$ . We denote  $pre$  and  $eff$  the functions which respectively associate preconditions and effects with each event:  $pre : \mathbb{E} \mapsto \mathcal{P}$ ,  $eff : \mathbb{E} \mapsto \mathcal{E}$ . Given the expression of  $E_{-1}$ , the application of  $eff$  to each element of the set is  $eff(ini_l) = l$  with  $l \in S_0$ , thus  $eff(E_{-1}) = S_0$ . Moreover, given a formula  $\psi \in \mathcal{P}$  and a partial state  $L$ ,  $L \models \psi$  is defined classically:  $L \models l$  if  $l \in L$ ,  $L \models \psi_1 \wedge \psi_2$  if  $L \models \psi_1$  and  $L \models \psi_2$ , and  $L \models \psi_1 \vee \psi_2$  if  $L \models \psi_1$  or  $L \models \psi_2$ .

Our work is a logical continuation of works such as [Hopkins and Pearl, 2007, Berreby et al., 2018, Batusov and Soutchanski, 2018], who attempted to link action languages and causation. To the best of our knowledge, Batusov and Soutchanski, and Berreby et al. are the first to give a definition of actual cause in action languages. However, each work has its own limitations that we try to address. In Batusov and Soutchanski’s paper, many working perspectives are mentioned [Batusov and Soutchanski, 2018]:

It is clear that a broader definition of actual cause requires more expressive action theories that can model not only sequences of actions, but can also include explicit time and concurrent actions. Only after that one can try to analyze some of the popular examples of actual causation formulated in philosophical literature. Some of those examples sound deceptively simple, but faithful modelling of them requires time, concurrency and natural actions.

At the moment, the proposed action language tackles both concurrency and time—at least discrete time. We will now introduce ‘natural actions’ that we denote exogenous events. The set  $\mathbb{E}$  is divided into two sub-sets:  $\mathbb{A}$ , which contains the actions carried out by an agent and thus subjected to a volition;  $\mathbb{U}$ , which contains the exogenous events—equivalent to :event in PDDL+ [Fox and Long, 2006] and triggered axioms in Event Calculus [Mueller, 2014]—which are triggered as soon as all the  $pre$  are fulfilled, therefore without the need for an agent to perform them. Thus, for exogenous events triggering conditions and preconditions are the same. In contrast, the triggering conditions for actions necessarily include preconditions but those are not sufficient. The triggering conditions of an action also include the volition of the agent or some kind of manipulation by another agent. To keep track of these subtleties that could be relevant in the causal inquiry we introduce triggering conditions represented as formulas of the language  $\mathcal{P}$ . We denote  $tri$  the function which associates triggering conditions with each event:  $tri : \mathbb{E} \mapsto \mathcal{P}$ .

Two events  $e, e' \in \mathbb{E}$  are said to be interfering if  $eff(e) \cup eff(e')$ —given the set notation introduced for elements of  $\mathcal{E}$ —is a set that is not coherent according to Definition 1.

**Definition 2** (context  $\kappa$ ). *Given an initial state  $S_0$ , the context denoted as  $\kappa$  is the octuple  $(\mathbb{E}, \mathbb{F}, pre, tri, eff, S_0, >, \mathbb{T})$ , where  $>$  is a partial order which represents priorities that ensure the primacy of one event over another when both are interfering.*

As mentioned earlier, effects indicate the changes to the fluents that are expected to happen if an event occurs. Because of the complexity of reality, it may turn out that causally the action has more or less effects than those attributed by  $\mathcal{E}$ . Let's take the example of an agent who wants to turn on a light by pressing a switch. In a first scenario, it is possible that the agent's action causes an overheating in the electrical circuit and triggers a fire. When formalising the action of switching on the light, it is not intuitive to take into account the overheating and then the fire as intrinsic effects. Besides affecting the generality of the formalisation, we previously mentioned that this could influence the ethical evaluation. In these cases, we will prefer to break down the process by introducing exogenous events. In the above fire example, we will therefore have an exogenous event corresponding to a fire outbreak—an agentless event—which will be triggered when a defective circuit is present and the switch is pressed. We are therefore in the presence of a causal chain. These cases where the action has more effects than those with which it has been formalised are typical cases where causality is necessary. In a second scenario, it may happen that the agent performs the action but that it does not have the expected effects simply because the light was already on. This does not prevent the action from having been performed, and we want to keep a trace of the event without having to consider that its effect has taken place. This is especially the case if the action has several effects and only one of them does not actually occur. This second case can be resumed as cases where some of the fluents of the state have already the value attributed by an effect. Since the effects that an event had at the time it occurred is a basic causal information on which we will rely—inextricably linked to imputability—it is important to keep track of them.

**Definition 3** (actual effects  $actualEff(E, L)$ ). *Given a context  $\kappa$ , the predicate  $actualEff(E, L)$  which associates a set of events  $E \in \mathbb{E}$  given a partial state  $L$ , to a partial state representing the actual effects of  $E$  when  $L$  is true, is defined as:*

$$\begin{aligned} actualEff(E, L) &= \bigcup_{e \in E} actualEff(\{e\}, L) \\ &= \{l_i, \exists e \in E, [\psi_i]l_i \in eff(e), L \models \psi_i, \text{ and } l_i \notin L\} \end{aligned}$$

For the sake of conciseness we adopt an update operator giving the resulting state when performing an event at a given state.

**Definition 4** (update operator  $\triangleright$ ). *Given a context  $\kappa$  and set of events  $E \in \mathbb{E}$ , the update operator which we use as follows  $S_t \triangleright E$  expresses  $S_t \setminus actualEff(E, S_t) \cup actualEff(E, S_t)$ .*

The information given by  $actualEff(E, L)$  and  $\triangleright$  can be equated to basic causal information given by the evolution of the world. Besides being causal, this information is directional since it is inconceivable in our semantics to say that the actual effect of the event is the cause of it. Therefore, we can rely on the events that occur and their actual effects to simulate the evolution of the world from the initial state  $S_0$ .

**Definition 5** (induced state sequence  $S_\kappa$ ). *Given a context  $\kappa$  and a sequence of events  $\epsilon = E_{-1}, E_0, \dots, E_n$ , such that  $n \leq |\mathbb{T}|$ , the induced state sequence of  $\epsilon$  is a sequence of complete states:  $S_\kappa(\epsilon) = S_0, S_1, \dots, S_{n+1}$  such that  $\forall t \in -1, \dots, n, S_{t+1} = S_t \triangleright E_t$ .*

Though this can be defined for every  $\epsilon$ , not all  $\epsilon$  are possible given (i) the need to satisfy preconditions, (ii) the concurrency of events that must respect priorities, and (iii) the triggering of events that must respect priorities too.

**Definition 6.** *Let  $\epsilon$  be a sequence of events  $\epsilon = E_{-1}, E_0, \dots, E_n$ , such that  $n \leq |\mathbb{T}|$ , and let's denote by  $S_\kappa(\epsilon) = S_0, \dots, S_{n+1}$  its induced state sequence. We shall say that  $\epsilon$  is:*

- *Executable in  $\kappa$ : if  $\forall t \in 0, \dots, n, S_t \models pre(E_t)$ .*
- *Concurrent correct with respect to  $\kappa$ :  $\neg \exists (e, e') \in E_t^2, e > e'$ .*
- *Trigger correct with respect to  $\kappa$ : if  $\forall t \in 0, \dots, n, \forall e' \in \mathbb{E}$  such that  $S_t \models tri(e')$ , then  $e' \in E_t$  or  $\exists e \in E_t, e > e'$ .*

- Valid in  $\kappa$ : if and only if, executable in  $\kappa$ , concurrent correct with respect to  $\kappa$ , and trigger correct with respect to  $\kappa$ .

Finally, if we consider only a set of timed actions as an input which we call scenario, we have:

**Definition 7** (traces  $\tau_{\sigma,\kappa}^e$  and  $\tau_{\sigma,\kappa}^s$ ). Given a scenario  $\sigma \subseteq \mathbb{A} \times \mathbb{T}$  and a context  $\kappa$ , the event trace  $\tau_{\sigma,\kappa}^e$  of  $\sigma, \kappa$  is the sequence of events  $\epsilon = E_{-1}, E_0, \dots, E_n$  valid in  $\kappa$ , such that:  $\forall t \in 0, \dots, n, \forall e \in E_t, e \in \mathbb{A} \Leftrightarrow (e, t) \in \sigma$ . Its induced state sequence is the state trace  $\tau_{\sigma,\kappa}^s$ .

**Example 3** (pollution—formalisation). Figure 2 illustrates how Examples 1 and 2 can be formalised in this action language. These examples are composed of two actions: the production of medicines and the production of connected speakers which we denote  $prod_m$  and  $prod_s$  respectively. The first is formalised as having two effects, the availability of  $k$  doses of medicine ( $m_k \in \mathbb{F}$ ) and the existence of wastewater ( $w_m \in \mathbb{F}$ ). The second also has two effects, employment for  $n$  individuals ( $e_n \in \mathbb{F}$ ) and the existence of wastewater ( $w_s \in \mathbb{F}$ ). We add to these actions two exogenous events: the fact of discharging wastewater ( $dis_w \in \mathbb{U}$ ) and the potable water plant fault ( $fau_p \in \mathbb{U}$ ). The event  $fau_p$  is simple, it is triggered when the pollution indicator of the lake is above the threshold ( $s_{sup} \in \mathbb{F}$ ) and it has as effect the damage of the inhabitants who are deprived of drinking water ( $d \in \mathbb{F}$ ). The event  $dis_w$  is triggered either when there is wastewater from the speakers factory ( $w_s$ ) or when there is wastewater from the medicine factory ( $w_m$ ) and  $w_m$  treatment plant is out of service ( $t_{os} \in \mathbb{F}$ ). This event raises the pollution indicator of the lake above the threshold ( $s_{sup}$ ).

$$\begin{aligned} pre(prod_m) &= \top, eff(prod_m) = m_k \wedge w_m \\ pre(prod_s) &= \top, eff(prod_s) = e_n \wedge w_s \\ pre(dis_w) &= w_s \vee (w_m \wedge t_{os}), eff(dis_w) = s_{sup} \\ pre(fau_p) &= s_{sup}, eff(fau_p) = d \end{aligned}$$




Figure 2: Diagram of the  $\kappa$  described in Example 3.

## 4 Adapted Causal Inquiry

Of the many fields studying causality, our approach is especially close to *tort law* whose interest is about causation in specific situations. Hence, works in this field are a good source of inspiration, with a large number of insights due to the still current discussions on the topic. In a series of influential papers [Wright, 1985, 1988], Wright demonstrates how essential a causal inquiry is in the process of determining *tort liability*—he emphasises the fundamental difference between causation and responsibility. In the process, he shows that this causal inquiry needs to be factual and independent of any policy choice. The argument we make is that such a causal inquiry is the same as the one that must be incorporated into proposals in computational ethics.

In the fable *The Wolf and the Lamb* [de La Fontaine, 1820], de La Fontaine makes us judges of the case confronting the two characters. In this story, the Lamb defends himself one by one against the increasingly absurd arguments put forward by the Wolf before being devoured. The latter begins by pleading that the Lamb disturbs his drink—he is accused of being the cause of his harm. Cleverly, the Lamb retorts that by ‘taking a drink of water in the stream more than twenty steps below him [the Wolf]’, it cannot be considered as the cause of the harm. Given this evidence invalidating the causal inquiry, the Wolf feels forced to ignore the factual aspects in order to plead his case. He abandons the causal inquiry

and turns to the subjectivity of responsibility. He then accuses the Lamb of having ‘said bad things’ of him, as if this made him guilty of the harm he is accused of. By doing this the Wolf choose to rely entirely on a subjective policy choice. The relevance of the argument is intuitively rejected by the reader because this conduct—even if accepted as possibly being inappropriate—has no causal relationship to the harm. The more the arguments follow one another and target the Lamb’s entourage, the more absurd they become as they lack causal support. Faced with the evidence of the Lamb who invalidates one argument after another by remaining factual, the Wolf loses his patience, ‘carries the lamb, and then eats him without any other why or wherefore’.

A satisfying liability analysis—which goal is to determine if a defendant is the ‘responsible cause’ of an injury—requires a factual and independent of policy choices causal inquiry. With the Wolf and the Lamb fable, de La Fontaine gives us an illustrative example of what Wright criticises in his papers: a process to determine responsibility for an injury in which the causal inquiry is flawed and polluted with subjective aspects—a process where causality and responsibility are conflated. Wright’s initial observation is that those two notions are too often conflated. The fact that ‘the phrase “*the cause*” is simply an elliptical way of saying “*the responsible cause*”’ [Wright, 1985] shows how thin the boundary between those notions is. To clarify this conflation, Wright describes the process to determine if an individual is legally responsible for an injury. This process has three stages: (i) *tortious-conduct inquiry*, where are identified the defendant’s conducts that could potentially imply legal responsibility (intentional, negligent, hazardous, . . .); (ii) *causal inquiry*, where is evaluated if the identified tortious conducts really contributed to cause the harm, i.e. if they can be considered as causes of the injury; (iii) *proximate-cause inquiry*, where other causes of the injury are considered, so as to evaluate if they mitigate or eliminate the defendant’s legal responsibility for the injury. Of those three stages, only the second is entirely factual and independent of policy choices. It determines if a conduct was a cause of the injury. The two others are subject to policy considerations that ‘determine which causes and consequences will give rise to liability’ [Wright, 1985]. Not to yield into the easy confusion between responsibility and causality, our goal is to propose a definition of actual causality suitable for a causal inquiry as presented by Wright, i.e. factual and independent of policy choices.

The actual causation definitions based solely on strong necessity—also known as counterfactual dependence—fail to capture the commonly accepted intuition on overdetermination cases (early preemption, late preemption, and symmetric overdetermination) [Hall and Paul, 2003, Menzies and Beebee, 2020]. The commonly used in law *But-for test* is one of those unsuccessful definitions—it is the basis of a significant part of the few works in computational ethics that integrates a mechanism to establish causal relationships [Lindner et al., 2017, Lindner and Bentzen, 2018]. This test states that ‘an act was a cause of an injury if and only if, but for the act, the injury would not have occurred’ [Wright, 1985]. ‘In the context of structural equations, this flawed account can be described as equating causation with counterfactual dependence’ [Beckers, 2021]. Let’s take Example 1 and apply the *But-for test* to it. Would the harm to the inhabitants of the village have occurred if the factory producing the speakers had not launched its production? Given the scenario the answer is yes, the harm would still have occurred due to the presence of  $w_m$  discharges and the state  $tp_{os}$  of  $w_m$  treatment plant. The production of speakers is therefore not a cause of the harm according to the *But-for test* because it is not necessary for the harm to occur. The same result is obtained if we apply it to the production of medicines. Hence, this test tells us that neither action is a cause of the harm—result which the reader will intuitively reject. Applying the *But-for test* to Example 2 gives exactly the same result. Given that overdetermination cases are not just hypothetical and rare cases (cases of pollution, suicide, economic loss, . . .), those strong necessity based approaches are not suitable for our purposes.

The dominant approach of actual causality—HP definition [Halpern, 2016]—deals with those cases, but at the cost of the factualness of the causal inquiry. This definition has the same roots than the *But-for test*, Hume’s definition of causation second formulation [Hume, 1748]:

we may define a cause to be an object followed by another, and where all objects, similar to the first, are followed by objects similar to the second. Or, in other words, where if the first object had not been, the second had never existed.

It is the result of an iterative process that originates in Pearl’s formalisation of Lewis’ vision [Lewis, 1973] in structural equations framework (SEF) [Pearl, 2000]. HP approach is more complete than the

But-for test in the sense that other elements in addition to counterfactual dependence were included in order to deal with some complex cases. One of those elements is interventionism. This assumption states that an event  $C$  causes a second event  $E$  if and only if, both events occur, and that, given an intervention allowing to fix the occurrence of a certain set of other events in the context—without being constrained to respect the physical coherence of the world—there is a context where if the first event had not occurred, the second would not have occurred either. This assumption is described by Beckers using SEF notation as [Beckers, 2021]:

**Interventionism** They all share the assumption [HP-style definitions] that the relation between counterfactual dependence and causation takes on the following form:  $C = c$  causes  $E = e$  iff  $E = e$  is counterfactually dependent on  $C = c$  given an intervention  $\vec{X} \leftarrow x$  that satisfies some conditions P. The divergence between these definitions is to be found in the condition P that should be satisfied.

Interventionism—that Beckers’ CNESS and Beckers and Vennekens’ BV definitions reject—introduces non factual elements to the causal inquiry which appear problematic even for Halpern [Halpern, 2018]:

if I fix BH [Billy hits] to zero here, I am sort of violating the way the world works. [...] I am contemplating counterfactuals are inconsistent with the equations but I seem to need to do that in order to get things to work out right. Believe me, we tried many other definitions.

In addition to non factual elements, the divergence on which ‘conditions P’ to apply can be equated with policy choices. These elements make HP-style definitions non adequate for our context.

The NESS test which subordinates necessity to sufficiency is an approach that deals with overdetermination cases [Wright, 1985, 1988, 2011, Baumgartner, 2013] and that satisfies our inquiry needs. Introduced by Wright in response to But-for test flaws, this test states that [Wright, 1985, 1988]:

A particular condition was a cause of a specific consequence if and only if it was a necessary element of a set of antecedent actual conditions that was sufficient for the occurrence of the consequence.

Unlike approaches mentioned above, it belongs to a second high impact approach family, regularity theories of causation [Andreas and Guenther, 2021]. Those theories are also based on Hume’s definition of causation, but on the first formulation. Specifically, the NESS test is closer to Mill’s interpretation of this formulation which introduced that there are potentially a multiplicity of distinct, but equally sufficient sets of conditions [Mill, 1843]. The NESS test is even closer to Mackie’s proposal. Indeed, unlike Mill’s vision whereby the cause is the sufficient set, Mackie considers that each element of the set is a cause [Mackie, 1980].

Examples 1 and 2 have both two potential sets of conditions sufficient to produce the inhabitants’ harm, each set related to one of the possible actions. In Example 2, the two sets have all their conditions being met, and both the production of speakers and medicines are a necessary element for the occurrence of one of the two sets. Thus, the NESS test indicates that these two actions are causes of the harm—as intuitively expected. In Example 1, only the set containing the production of speakers has all its conditions met. This action is therefore considered by the NESS test to be a cause of the harm, and we will say that it is a cause that preempts the effects of the production of medicines. The production of medicines will therefore not be considered a cause—as intuitively expected.

The actual causation definition we propose is an action languages suitable formalisation of Wright’s NESS test. Even if accepted by influential counterfactual theories of causation authors as embodying our basic intuition of causation—such as Pearl [Beckers, 2021]—criticism of the use of logic as formalism has prevented the popularisation of this test. What is argued is the inadequacy of logical sufficiency and logical necessity to formalise these intuitions. Recent works have shown that rejecting the formalism is not a reason to reject the idea behind it by successfully formalising the NESS test in causal calculus [Bochman, 2018a] and in the structural equations framework [Beckers, 2021]. It is conceivable to work on a way of compiling existing action languages problems and translating them into SEF. However, works have shown SEF flaws [Bochman, 2018b] and that in complex evolving contexts [Hopkins and Pearl,

2007, Batusov and Soutchanski, 2018] like ours, this translating approach is not necessarily desirable [Hopkins and Pearl, 2007]:

Structural causal models are excellent tools for many types of causality-related questions. Nevertheless, their limited expressivity render them less than ideal for some of the more delicate causal queries, like actual causation. These queries require a language that is suited for dealing with complex, dynamically changing situations.

Now that we have defined the formalism in which we will represent the problems and the desired characteristics of the causal inquiry, all that remains to link automated planning and causality—and computational ethics and causality by the same occasion—is to introduce our action languages suitable definition of causation. Since we are in an operational framework, we can take some distance from metaphysical considerations and assume ‘that [causal laws] they are deterministic, and that they permit neither backwards causation nor causation across a temporal gap’ [Hall and Paul, 2003].

## 5 Actual Causality

In the context of action languages, we consider that a first event is an actual cause of a second event if and only if the occurrence of the first is a NESS-cause of the triggering of the second. As commonly accepted by philosophers, the relation of causality we aim to define links two events. However, ‘events are not the only things that can cause or be caused’ [Lewis, 1973]. Action languages represent the evolution of the world as a succession of states produced by the occurrence of events, thus introducing states between events. Therefore, we need to define causal relations where the cause are events and the effect are formulas of the language  $\mathcal{P}$ . This section will introduce definitions which establish such a relation based on Wright’s NESS test of causation.

**Definition 8** (causal setting  $\chi$ ). *The action language causal setting denoted  $\chi$  is the couple  $(\sigma, \kappa)$  with  $\sigma$  a scenario and  $\kappa$  a context.*

From now on, when reference is made to events and states, they will be those from  $\tau_{\sigma, \kappa}^e$  and  $\tau_{\sigma, \kappa}^s$  respectively. Thus, the set of all events which actually occurred at time point  $t$  is  $E^x(t) = \tau_{\sigma, \kappa}^e(t)$ . Following the same reasoning, the actual state at time point  $t$  is  $S^x(t) = \tau_{\sigma, \kappa}^s(t)$ .

**Definition 9** (direct NESS-causes). *Given a causal setting  $\chi$  and a partial state  $W \subseteq Lit_{\mathbb{F}}$  that we call backing, the occurrence of events of the set  $C = \{(e, t), e \in E^x(t), t \in \mathbb{T}\}$  are direct NESS-causes—backed by  $W$ —of the truthfulness of the formula  $\psi$  at  $t_\psi$ , denoted  $C \underset{W}{\rightsquigarrow} \psi, t_\psi$ , iff:*

- *Causal sufficiency and minimality of  $W$ :  $W \models \psi$  and  $\forall W' \subset W, W' \not\models \psi$ .*  
*There is a decreasing sequence  $t_1, \dots, t_k$  and a partition  $W_1, \dots, W_k$  of  $W$  such that  $\forall i \in 1, \dots, k$ , given  $C(t_i) = C \cap E^x(t_i)$ :*
- *Weak necessity and minimality of  $C$ :  $S^x(t_i) \triangleright C(t_i) \models W_i$  and  $\forall C' \subset C(t_i), S^x(t_i) \triangleright C' \not\models W_i$ .*
- *Persistence of necessity:  $\forall t, t_i < t \leq t_\psi, S^x(t) \models W_i$ .*

Wright’s NESS test is based on three main principles which are formalised in Definition 9: (i) sufficiency of a set, (ii) weak necessity of the conditions in that set, and (iii) actuality of the conditions. (i) In this definition, the sufficient set is the partial state  $W$ . More precisely, given the directionality embedded in Section 3 semantics, we have causal sufficiency that Wright differentiates from logical sufficiency [Wright, 2011]: ‘The successional nature of causation is incorporated in the concept of causal sufficiency, which is defined as the complete instantiation of all the conditions in the antecedent of the relevant causal law’. Moreover, this definition introduces the constraint of necessity and sufficiency minimality which has been proven to be essential for regularity theories of causation [Wright, 2011, Baumgartner, 2013, Andreas and Guenther, 2021]. (ii) Definition 9 formalises weak necessity by subordinating necessity to sufficiency achieving that [Wright, 2011]: ‘a causally relevant factor need merely be necessary for the sufficiency of a set of conditions sufficient for the occurrence of the consequence, rather than being necessary for the consequence itself’. It is worth mentioning that the condition  $S^x(t_i) \not\models W_i$ —intuitively

expected when referring to necessity—is included in the minimality condition by the case where  $C' = \emptyset$ , thus  $S^x(t_i) = S^x(t_i) \triangleright \emptyset \not\equiv W_i$ . (iii) The actuality of the conditions is assured by the use of actual occurrence of events, which is implied by the presence of  $E^x(t_i)$  in  $C(t_i) = C \cap E^x(t_i)$ .

Once causal sufficiency and minimality of the partial state  $W$  is defined, the causal inquiry is conducted by a recursive reasoning on a partition of  $W$ . The goal of this recursive reasoning is to identify the events which occurrence was necessary to the sufficiency of  $W$ . This reasoning is done by going back in time and analysing the information given by  $\tau_{\sigma,\kappa}^s(t)$  and  $\tau_{\sigma,\kappa}^e(t)$ . Two limit cases can be identified. The first is when  $W_k$  is empty before its corresponding time  $t_k$  is equivalent to  $t = 0$ , meaning that all occurrences of events necessary for the sufficiency of  $W$  have been identified. When this is not the case, it means that there are fluent literals in  $W$  that were true in the initial state  $S^x(0)$  and which value has not changed until  $S^x(t_\psi)$ . In this second case, the set  $C$  will contain the events  $ini_l \in E^x(-1)$  whose  $l$  remains in  $W_k$ —events which symbolise events in the past beyond the framework of formalisation.

In practice, it is possible to study what will be considered as the direct NESS-causes of the truthfulness of  $\psi$  at  $t_\psi$  for each form that  $\psi$  may take. In the case where  $\psi$  is a fluent literal  $l$ , the NESS-causes will be the last occurrences of events to have made  $l$  true before or at  $t_\psi$ . In this basic case  $W$  is the singleton which unique element is that literal. This basic causal information is the one embedded in Section 3 action language semantics. In the case where  $\psi$  is a conjunction  $\psi = l_1 \wedge \dots \wedge l_m$  of fluent literals, the NESS-causes will be all the occurrence of events that are NESS-causes of the truthfulness of one of the literals  $l_i$  in the conjunction at  $t_\psi$ . Finally, the case where  $\psi$  is a disjunction, and more generally a disjunctive normal form, is by far the more interesting and challenging. Indeed, it is in this case that we can be confronted to situations of overdetermination. Whenever  $\psi$  is a disjunction, this means that there is a minimal causal sufficient backing  $W$  for each disjunct. Each of these backings is a possible way to cause the truthfulness of the formula  $\psi$  at  $t_\psi$ —in the same spirit as Beckers' paths [Beckers, 2021]. Example 4 illustrates how Definition 9 handles one of those challenging situations.




Figure 3: Electrical circuit consisting of a voltage source, three switches, and an individual connected to electrodes.




Figure 4: Evolution of fluents given  $\kappa$  in Example 4.

**Example 4** (parallel switches and Milgram). Consider Figure 3 simple electric circuit inspired by Milgram's experiment [Milgram, 1963]. This circuit is made up of a voltage source, an individual strapped and connected to electrodes, and three switches connected in parallel. The positive literals  $l_1, l_2, l_3, l_4 \in Lit_{\mathbb{F}}$  represent the closed state of each switch and the voltage source respectively—their respective complement thus represents the opened state.  $\psi = (l_1 \wedge l_4) \vee (l_2 \wedge l_4) \vee (l_3 \wedge l_4)$  where  $\psi \in \mathcal{P}$  represents the triggering conditions for the strapped individual being electrocuted. Thus, three backings are possible to cause  $\psi$ :  $W = \{l_1, l_4\}$ ,  $W' = \{l_2, l_4\}$ , and  $W'' = \{l_3, l_4\}$ .  $e_1, e_2, e_3 \in \mathbb{E}$  are the events which intrinsic effect is to close each switch respectively,  $e_4 \in \mathbb{E}$  is an event which intrinsic effect is to close the voltage source, and  $e_{-1} \in \mathbb{E}$  is the event which intrinsic effect is to open the first switch. We assume that the situation involves five agents: the one strapped and four others—each controlling one of the four components of the circuit. The studied sequence illustrated by Figure 4 and given by  $\tau_{\sigma,\kappa}^e$  and  $\tau_{\sigma,\kappa}^s$  for this case is:

$$E^x(-1) = \{ini_{l_1}, ini_{\bar{l}_2}, ini_{\bar{l}_3}, ini_{\bar{l}_4}\}$$

$$S^x(0) = \{l_1, \bar{l}_2, \bar{l}_3, \bar{l}_4\}, E^x(0) = \{e_{-1}, e_2\}$$

$$S^X(1) = \{\bar{l}_1, l_2, \bar{l}_3, \bar{l}_4\}, E^X(1) = \{e_3, e_4\}$$

$$S^X(2) = \{\bar{l}_1, l_2, l_3, l_4\}$$

Given the above traces,  $\psi$  is true at  $t = 2$  by both  $W'$  and  $W''$ .

The question that arises in Example 4 is: what are the causes of  $\psi$  being true at  $t = 2$ ? Said in another way, what are the causes of the strapped individual being electrocuted at  $t = 2$ ? Batusov and Soutchanski's proposal will consider  $ini_{l_1}$  and  $e_4$  as 'achievement causes', and  $e_2$  as a 'maintenance cause'—'causes responsible for protecting a previously achieved effect, despite potential threats that could destroy the effect' [Batusov and Soutchanski, 2018]—this given that we do not consider  $e_3$  because it occurs at the same time as  $e_4$ . Considering factuality as an essential feature of a causal inquiry, the presence of  $ini_{l_1}$  in the causes seems unacceptable. Factually,  $ini_{l_1}$  plays no role in the truthfulness of  $\psi$  at  $t = 2$ . Definition 9 gives the sets  $\{e_2, e_4\}$  and  $\{e_3, e_4\}$  which union gives the answer  $\{e_2, e_3, e_4\}$ .

The interpretation given by Batusov and Soutchanski of Example 4 is not the only possible divergent interpretation. We wondered whether answer  $\{e_2, e_4\}$  was not more satisfactory given that, even if both  $l_2$  and  $l_3$  are true at  $t = 2$ , the precedence of  $l_2$  should be taken into account. However, this intuition appears as conflating causality and responsibility. If we strictly limit ourselves to a factual causal inquiry as prescribed by Wright, both  $e_2$  and  $e_3$  are causes of the truthfulness of  $\psi$  at  $t = 2$ . The intuition leading us to take into account the precedence of  $e_2$  belongs to Wright's *proximate-cause inquiry* and not to the causal inquiry. Indeed, once  $e_2$  and  $e_3$  are identified as causes, there is a policy choice for which the precedence of  $e_2$  mitigates or eliminates the responsibility of  $e_3$  for the final effect. We suspect that Batusov and Soutchanski's choices—which led them to consider  $ini_{l_1}$  as a cause—were influenced by this same intuition, but taken even further.

Definition 9 gives us essential information about causal relations by looking to the actual effects of events. However, the set of direct NESS-causes of an effect may include exogenous events that are not necessarily relevant. This is especially true in a framework such as ours, where we are interested in the ethical dimension of an agent's decisions—thus actions. It is therefore essential to establish a causal chain by going back in time in order to find the set of actions that led to the effect. To this end, we must broaden our vision to look not only at the actual effects of events which are direct NESS-causes, but also (i) at the events that caused those events to be triggered and (ii) at the events that caused those events to have their actual effects.

**Example 5** (causing events to have their actual effects). *Consider the literals  $l_1, l_2, l_3, l_{c_1}, l_{c_3} \in Lit_{\mathbb{F}}$ , the formula  $\psi = l_1 \wedge l_2 \wedge l_3$ , events  $e, e' \in \mathbb{E}$  where there respective effects are  $eff(e) = \{[l_{c_1}]\bar{l}_1, [\top]l_2, [l_{c_3}]l_3\}$  and  $eff(e') = \{[\top]\bar{l}_{c_1}, [\top]l_{c_3}\}$ , and the sequence given by  $\tau_{\sigma, \kappa}^e$  and  $\tau_{\sigma, \kappa}^{e'}$ :*

$$E^X(-1) = \{ini_{l_1}, ini_{\bar{l}_2}, ini_{\bar{l}_3}, ini_{\bar{l}_4}, ini_{l_{c_1}}, ini_{\bar{l}_{c_3}}\}$$

$$S^X(0) = \{l_1, \bar{l}_2, \bar{l}_3, l_{c_1}, \bar{l}_{c_3}\}, E^X(0) = \{e'\}$$

$$S^X(1) = \{l_1, \bar{l}_2, \bar{l}_3, \bar{l}_{c_1}, l_{c_3}\}, E^X(1) = \{e\}$$

$$S^X(2) = \{l_1, l_2, l_3, \bar{l}_{c_1}, l_{c_3}\}$$

Given the above sequence, Definition 9 gives us the direct NESS-cause relation  $C \xrightarrow[W]{\rightsquigarrow} \psi, t_\psi$  where  $C$  is the set  $\{(e, 1), (ini_{l_1}, -1)\}$ .

In Example 5, the actual effects of the occurrence of  $e$  were  $actualEff(\{e\}, S^X(1)) = \{l_2, l_3\}$ . In order to determine the desired causal chain, one of the steps requires to ask ourselves what events caused  $e$  to have those effects—inquiry concerning exclusively conditional effects which condition is not  $[\top]$ —we distinguish two cases. In our example, two effects are concerned,  $[l_{c_1}]\bar{l}_1$  and  $[l_{c_3}]l_3$ , each one representing a case. The effect  $[l_{c_1}]\bar{l}_1$  corresponds to the case where the complement of the condition  $[l_{c_1}]$  has been direct NESS-caused, thus causing  $e$  to 'maintain'  $l_1$ . The effect  $[l_{c_3}]l_3$  corresponds to the case where the condition  $[l_{c_3}]$  has been direct NESS-caused, thus causing  $e$  to 'produce'  $l_3$  as an actual effect. The predicate *after*  $(E, L_p, L_m)$ —inspired by Khan and Lespérance's work [Khan and

Lespérance, 2021]—gives the formula to direct NESS-cause in order to be considered a cause of an event having its actual effects. In the discussed example this formula is  $\psi' = \overline{l_{c_1}} \wedge l_{c_3}$ .

**Definition 10** ( $after(E, L_p, L_m)$ ). Given a causal setting  $\chi$ , a set of events  $E \in E^X(t)$ , and partial states  $L_m, L_p, W_{\psi'} \subseteq Lit_{\mathbb{F}}$  such that  $S^X(t) \models L_m$  and  $S^X(t) \not\models L_p$ , the predicate  $after(E, L_p, L_m) = \psi'$  with  $\psi' = \bigwedge_{l \in W_{\psi'}} l$  such that:

- *Necessity and minimality of E*:  $W_{\psi} \triangleright E \models L_p \cup L_m$  and  $\forall E' \subset E, W_{\psi} \triangleright E' \not\models L_p \cup L_m$ .
- *Monotonicity*:  $\forall W', W_{\psi} \subseteq W', W' \triangleright E \models L_p \cup L_m$ .

Having introduced the predicate  $after(E, L_p, L_m)$ , we can now introduce NESS-causes that are found relying on the establishment of the causal chain.

**Definition 11** (NESS-causes). Given a causal setting  $\chi$ , the direct NESS-cause relation  $C \xrightarrow[W]{\sim} \psi, t_{\psi}$ , and the decreasing sequence  $t_1, \dots, t_k$  induced by the existing partition  $W_1, \dots, W_k$  of the backing  $W$ , the occurrence of events of the set  $C' = \{(e, t), e \in E^X(t), t \in \mathbb{T}\}$  are NESS-causes of the truthfulness of the formula  $\psi$  at  $t_{\psi}$  iff:

- *Base case*:  $C' = C$ .
- *Recursive case*: Given the sets  $C_R = C \setminus C'$  and  $C_O = C' \setminus C$  of ‘remaining’ and ‘overwritten’ occurrence of events respectively, and the partitions of  $C$  and  $C_R$  matching the decreasing sequence  $t_1, \dots, t_k, C(t_1), \dots, C(t_k)$  and  $C_R(t_1), \dots, C_R(t_k)$  respectively, there is a partition of  $C_O = \bigcup_{i \in 0, \dots, k} C_{O_i}$  (not necessarily monotonic in time) such that:

$$\forall i, C_R(t_i) \neq \emptyset \text{ and } (e, t) \in C_{O_i} \text{ are NESS-causes of } \psi' \text{ true at } t_i, \text{ where } \psi' = tri(C_R(t_i)) \wedge after(C_R(t_i), W_i \cap actEff(C_R(t_i), S^X(t_i)), W_i \setminus actEff(C_R(t_i), S^X(t_i)) \cup W_{i+1} \cup \dots \cup W_k)$$

Then, the set of NESS-causes  $D = C \setminus C_R \cup C'$  is called a set of decisional causes if  $D \subseteq \mathbb{A}$ .

Having determined the causal relations linking events and formulas of the language, we can now give a suitable for action languages definition of actual causality.

**Definition 12** (actual cause). Given a causal setting  $\chi$ , the actual causes of the event  $e \in E^X(t_{\psi})$  are the NESS-causes of  $tri(e)$ , i.e. the truthfulness of the triggering conditions of  $e$  at  $t_{\psi}$ .

**Example 6** (pollution—causal relations). Figure 5 illustrates Example 2 sequence of events as well as the causal relations that can be established based on the proposed definitions. We can see that as expected, both the launching of the production of connected speakers and medicines are NESS-causes of the harm, and that these two actions are also actual causes of the potable water plant fault. The reader may verify that given the sequence of events in Example 1, the only action identified as a cause of the harm is the launching of the production of connected speakers—as expected.




Figure 5: Causal relations in Example 2.

## 6 Conclusion

The contribution of this paper is double. First, we link automated planning and causality by continuing the momentum established by recent papers [Batusov and Soutchanski, 2018, Berreby et al., 2018] of proposing an action languages suitable definition of actual causality. Second, we link computational ethics and causality by showing the necessity of a mechanism allowing to establish complex causal relations in the simulation of ethical reasoning—a practice that is still rare in the domain—and by proposing such a mechanism. Through our actual causation definition proposal we address two of what we consider the main remaining limitations of the linking automated planning and causality venture. First, not to yield into the easy confusion between responsibility and causality, our proposal is suitable for a factual and independent of policy choices causal inquiry. Second, not to disregard the much debated cases of overdetermination, our proposal is based on an action language semantics allowing concurrency of events. By taking as a basis Wright’s NESS test we are able to manage these cases satisfactorily. To the best of our knowledge, no other action languages suitable definition of actual causality has been able to handle those complex cases, yet essential. Our approach thus allows to handle complex cases of causality in decision making applications. Integrating this approach into an ethical decision making framework allows us by the same occasion to contribute to the progress of computational ethics by enabling us to deal with situations that were previously out of reach.

In future work we intend to extend the definition of causality by including the relation ‘prevent’ [Berreby et al., 2018]. In Wright’s conception of causality, causality can only be sufficient if we take into account—in addition to the positive causes—the conditions that were not true and whose absence was a necessary condition for the occurrence of the result. Then, the events being causes of their absence are also causes of the result. By working on fluent literals, our definition of causation already takes this notion into account. If we extend this reasoning, we could also take the case where the result did not occur because one of these negative conditions was made true. In such a case, the events being causes of the negative condition are causes of the non-occurrence of the result. This relation seems to be indispensable in ethical reasoning. In the example used throughout this paper, if an agent had the possibility to prevent the harm, the fact that he did not is a relevant information for the ethical evaluation. Thus, we intend to define this causal relation given our more complex framework with events concurrency and disjunction.

Furthermore, considering the implications for ethical evaluation, we feel essential to investigate in future work the difference between the concepts of ‘enables’ and ‘cause’—a point on which some interesting leads from various perspectives have been given [Martin, 1994, Lewis, 1997, Sloman et al., 2009, Berreby et al., 2018, Choi and Fara, 2021]. Consider that the precondition for the action burglarising a house is that the door was left open. If an agent decides to go and rob the house, it seems correct to say that another agent who forgot to close the door has enabled the robbery, not that he is an actual cause of it. Indeed, in this case the event of interest is an action of an agent with a volition. The ‘enables’ relation seems to be definable by replacing the *tri* function by *pre* in the actual cause definition. The possible nuances of causality in relationship with the ‘enabling’ notion do not seem to have been investigated by analytical philosophy in depth, considering all possible theoretical implications. However, applied bioethics seems to be an exciting theoretical ground in which the difference between causing and enabling an event appears crucial and even well conceptualized. This is particularly true when we deal with end-of-life issues such as euthanasia and assisted suicide. As numerous bioethics committees have shown, the crucial distinction between killing and letting die is based upon the assumption of different kind of moral responsibility when causing or allowing an event to happen. Thus, we intend to start from these issues to open a path of research that could definitively enrich our causality theory.

## References

- Holger Andreas and Mario Guenther. Regularity and Inferential Theories of Causation. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2021 edition, 2021. URL <https://plato.stanford.edu/archives/fall12021/entries/causation-regularity/>.
- United Nations General Assembly. Universal Declaration of Human Rights, December 1948. URL [https://www.un.org/en/ga/search/view\\_doc.asp?symbol=A/RES/217\(III\)](https://www.un.org/en/ga/search/view_doc.asp?symbol=A/RES/217(III)). resolution 217 A (III), A/RES/3/217 A.

- Vitaliy Batusov and Mikhail Soutchanski. Situation Calculus Semantics for Actual Causality. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 1744–1752, New Orleans, Louisiana, USA, 2018. AAAI Press. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16891>.
- Michael Baumgartner. A Regularity Theoretic Approach to Actual Causation. *Erkenntnis*, 78(1):85–109, December 2013. ISSN 1572-8420. doi: 10.1007/s10670-013-9438-3. URL <https://doi.org/10.1007/s10670-013-9438-3>.
- Sander Beckers. The Counterfactual NESS Definition of Causation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(7):6210–6217, May 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16772>.
- Sander Beckers and Joost Vennekens. A principled approach to defining actual causation. *Synth.*, 195(2):835–862, 2018. doi: 10.1007/s11229-016-1247-1. URL <https://doi.org/10.1007/s11229-016-1247-1>.
- Helen Beebe, Christopher Hitchcock, and Peter Menzies. *The Oxford Handbook of Causation*. OUP Oxford, November 2009. ISBN 978-0-19-927973-9. URL <https://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199279739.001.0001/oxfordhb-9780199279739>.
- Fiona Berreby, Gauvain Bourgne, and Jean-Gabriel Ganascia. Modelling Moral Reasoning and Ethical Responsibility with Logic Programming. In *20th International Conference, Logic for Programming, Artificial Intelligence, and Reasoning*, volume 9450 of *Lecture Notes in Computer Science*, pages 532–548, Suva, Fiji, November 2015. Springer. doi: 10.1007/978-3-662-48899-7\_37. URL <https://hal.sorbonne-universite.fr/hal-01395030>.
- Fiona Berreby, Gauvain Bourgne, and Jean-Gabriel Ganascia. A Declarative Modular Framework for Representing and Applying Ethical Principles. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS '17*, pages 96–104, Richland, SC, May 2017. International Foundation for Autonomous Agents and Multiagent Systems. URL <https://hal.archives-ouvertes.fr/hal-01564675>.
- Fiona Berreby, Gauvain Bourgne, and Jean-Gabriel Ganascia. Event-Based and Scenario-Based Causality for Computational Ethics. In *AAMAS 2018 - 17th International Conference on Autonomous Agents and Multiagent Systems*, pages 147–155, Stockholm, Sweden, July 2018. International Foundation for Autonomous Agents and Multiagent Systems. URL <https://hal.archives-ouvertes.fr/hal-01982090>.
- Alexander Bochman. Actual Causality in a Logical Setting. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018*, pages 1730–1736, Stockholm, Sweden, 2018a. ijcai.org. doi: 10.24963/ijcai.2018/239. URL <https://doi.org/10.24963/ijcai.2018/239>.
- Alexander Bochman. On Laws and Counterfactuals in Causal Reasoning. In Michael Thielscher, Francesca Toni, and Frank Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference*, pages 494–503, Tempe, Arizona, 2018b. AAAI Press. URL <https://aaai.org/ocs/index.php/KR/KR18/paper/view/17990>.
- Vincent Bonnemains, Claire Saurel, and Catherine Tessier. Embedded ethics: some technical and ethical challenges. *Ethics and Information Technology*, 20(1):41–58, March 2018. ISSN 1572-8439. doi: 10.1007/s10676-018-9444-x. URL <https://doi.org/10.1007/s10676-018-9444-x>.
- Gauvain Bourgne, Camilo Sarmiento, and Jean-Gabriel Ganascia. ACE modular framework for computational ethics: dealing with multiple actions, concurrency and omission. In *1st International Workshop on Computational Machine Ethics*, Online event, 2021. CEUR-WS.org. URL <http://www.cse.unsw.edu.au/~cme2021/>.
- Sungho Choi and Michael Fara. Dispositions. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2021 edition, 2021. URL <https://plato.stanford.edu/archives/spr2021/entries/dispositions/>.
- Thomas d’Aquin. *Summa theologica (Somme théologique)*, volume Secunda secundae : la morale particulière, question 64 : L’homicide. s.n., 1266. URL [http://docteurangelique.free.fr/saint\\_thomas\\_d\\_aquin/oeuvres\\_completes.html](http://docteurangelique.free.fr/saint_thomas_d_aquin/oeuvres_completes.html).

- Jean de La Fontaine. *Fables de La Fontaine*, volume vol. 2 of *Fables de La Fontaine*. Méquignon-Marvis, 1820. URL <https://aestheticrealism.net/poems/the-wolf-and-the-lamb-by-jean-de-la-fontaine/>.
- Philippa Foot. The Problem of Abortion and the Doctrine of the Double Effect. *Oxford Review*, 5:5–15, 1967. URL <https://philarchive.org/rec/F00TP0-2>.
- Maria Fox and Derek Long. Modelling Mixed Discrete-Continuous Domains for Planning. *Journal of Artificial Intelligence Research*, 27:235–297, October 2006. ISSN 1076-9757. doi: 10.1613/jair.2044. URL <https://jair.org/index.php/jair/article/view/10471>.
- Jean-Gabriel Ganascia. Non-monotonic Resolution of Conflicts for Ethical Reasoning. In Robert Trappl, editor, *A Construction Manual for Robots' Ethical Systems: Requirements, Methods, Implementations*, Cognitive Technologies, pages 101–118. Springer International Publishing, Cham, 2015. ISBN 978-3-319-21548-8. doi: 10.1007/978-3-319-21548-8\_6. URL [https://doi.org/10.1007/978-3-319-21548-8\\_6](https://doi.org/10.1007/978-3-319-21548-8_6).
- Malik Ghallab, Craig Knoblock, David Wilkins, Anthony Barrett, Dave Christianson, Marc Friedman, Chung Kwok, Keith Golden, Scott Penberthy, David Smith, Ying Sun, and Daniel Weld. PDDL - The Planning Domain Definition Language. *Technical Report CVC TR-98-003/DCS TR-1165*, Yale Center for Computational Vision and Control, August 1998. URL [https://www.researchgate.net/publication/2278933\\_PDDL\\_-\\_The\\_Planning\\_Domain\\_Definition\\_Language](https://www.researchgate.net/publication/2278933_PDDL_-_The_Planning_Domain_Definition_Language).
- Naveen Sundar Govindarajulu and Selmer Bringsjord. On automating the doctrine of double effect. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI'17*, pages 4722–4730, Melbourne, Australia, August 2017. AAAI Press. ISBN 978-0-9992411-0-3. URL <https://doi.org/10.24963/ijcai.2017/658>.
- Ned Hall and Laurie Ann Paul. Causation and Pre-emption. In Peter Clark and Katherine Hawley, editors, *Philosophy of Science Today*. Oxford University Press, Oxford, New York, May 2003. ISBN 978-0-19-925055-4. URL <https://global.oup.com/academic/product/philosophy-of-science-today-9780199250554?cc=fr&lang=en&#>.
- Joseph Y. Halpern. *Actual Causality*. The MIT Press, 2016. ISBN 978-0-262-03502-6. URL <https://www.jstor.org/stable/j.ctt1f5g5p9>.
- Joseph Y. Halpern. Actual Causality: A Survey: Joseph Halpern, June 2018. URL <https://www.youtube.com/watch?v=hXnCX2pJ0sg>. 26:38-27:21.
- Patrik Haslum, Nir Lipovetzky, Daniele Magazzeni, and Christian Muise. *An Introduction to the Planning Domain Definition Language*. Number 42 in Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, April 2019. ISBN 978-1-68173-512-2.
- Mark Hopkins and Judea Pearl. Causality and Counterfactuals in the Situation Calculus. *Journal of Logic and Computation*, 17(5):939–953, October 2007. ISSN 0955-792X. doi: 10.1093/logcom/exm048. URL <https://doi.org/10.1093/logcom/exm048>.
- David Hume. *Enquête sur l'entendement humain*. Number 1305 in GF. Flammarion, Paris, 2006 edition, 1748. ISBN 978-2-08-071305-6 978-2-08-025679-9. URL <https://catalogue.bnf.fr/ark:/12148/cb40216803q>.
- Emmanuel Kant. *Fondements de la métaphysique des moeurs*. Philosophie. Hachette Livre BNF, 1915 edition, May 2013. ISBN 978-2-01-280275-9.
- Shakil M. Khan and Yves Lespérance. Knowing Why: On the Dynamics of Knowledge about Actual Causes in the Situation Calculus. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '21*, pages 701–709, Richland, SC, May 2021. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 978-1-4503-8307-3. URL <https://www.ifaamas.org/Proceedings/aamas2021/pdfs/p701.pdf>.
- David Lewis. Causation. *Journal of Philosophy*, 70(17):556–567, 1973. doi: 10.2307/2025310. URL <https://philpapers.org/rec/LEWC>. Publisher: Oxford Up.
- David Lewis. Finkish Dispositions. *The Philosophical Quarterly*, 47(187):143–158, April 1997. ISSN 0031-8094. doi: 10.1111/1467-9213.00052. URL <https://doi.org/10.1111/1467-9213.00052>.

- Raynaldio Limarga, Maurice Pagnucco, Yang Song, and Abhaya Nayak. Non-monotonic reasoning for machine ethics with situation calculus. In *AI 2020: Advances in Artificial Intelligence: 33rd Australasian Joint Conference, AI 2020 Canberra, ACT, Australia, November 29–30, 2020, Proceedings*, pages 203–215, Canberra, ACT, Australia, 2020. Springer, Springer Nature. doi: 10.1007/978-3-030-64984-5\_16. URL <https://www.springerprofessional.de/en/non-monotonic-reasoning-for-machine-ethics-with-situation-calcul/18630242>.
- Felix Lindner and Martin Mose Bentzen. A Formalization of Kant’s Second Formulation of the Categorical Imperative. In *14th International Conference on Deontic Logic and Normative Systems*, Utrecht, the Netherlands, 2018. URL <https://orbit.dtu.dk/en/publications/a-formalization-of-kants-second-formulation-of-the-categorical-im>.
- Felix Lindner, Martin Mose Bentzen, and Bernhard Nebel. The HERA approach to morally competent robots. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6991–6997, Vancouver, Canada, September 2017. doi: 10.1109/IROS.2017.8206625. URL <https://ieeexplore.ieee.org/document/8206625>. ISSN: 2153-0866.
- John Leslie Mackie. *The Cement of the Universe: A Study of Causation*. Clarendon Library of Logic and Philosophy. Oxford University Press, Oxford, 1980. ISBN 978-0-19-824642-8. doi: 10.1093/0198246420.001.0001. URL <https://oxford.universitypressscholarship.com/10.1093/0198246420.001.0001/acprof-9780198246428>.
- C. B. Martin. Dispositions and Conditionals. *The Philosophical Quarterly (1950-)*, 44(174):1–8, 1994. ISSN 0031-8094. doi: 10.2307/2220143. URL <https://www.jstor.org/stable/2220143>. Publisher: [Oxford University Press, University of St. Andrews, Scots Philosophical Association].
- Peter Menzies and Helen Beebe. Counterfactual Theories of Causation. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2020 edition, 2020. URL <https://plato.stanford.edu/archives/win2020/entries/causation-counterfactual/>.
- Stanley Milgram. Behavioral Study of obedience. *The Journal of Abnormal and Social Psychology*, 67(4):371–378, 1963. ISSN 0096-851X. doi: 10.1037/h0040525. URL <https://psycnet.apa.org/record/1964-03472-001>. Place: US Publisher: American Psychological Association.
- John Stuart Mill. *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence, and the Methods of Scientific Investigation*, volume 1 of *Cambridge Library Collection - Philosophy*. Cambridge University Press, Cambridge, 2011 edition, 1843. ISBN 978-1-108-04088-4. doi: 10.1017/CBO9781139149839. URL <https://www.cambridge.org/core/books/system-of-logic-ratiocinative-and-inductive/290C43FBA4DC7022540D58E7EC49B1C2>.
- Erik T. Mueller. *Commonsense Reasoning: An Event Calculus Based Approach*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2 edition, 2014. ISBN 978-0-12-801647-3.
- Sven Nyholm and Jilles Smids. The Ethics of Accident-Algorithms for Self-Driving Cars: an Applied Trolley Problem? *Ethical Theory and Moral Practice*, 19(5):1275–1289, November 2016. ISSN 1572-8447. doi: 10.1007/s10677-016-9745-2. URL <https://link.springer.com/article/10.1007/s10677-016-9745-2>.
- Judea Pearl. *Causality: models, reasoning, and inference*. Cambridge University Press, Cambridge, U.K. ; New York, 2000. ISBN 978-0-521-89560-6 978-0-521-77362-1. URL <https://psycnet.apa.org/record/2000-07461-000>.
- Edwin P. D. Pednault. ADL: exploring the middle ground between STRIPS and the situation calculus. In *Proceedings of the first international conference on Principles of knowledge representation and reasoning*, pages 324–332, San Francisco, CA, USA, December 1989. Morgan Kaufmann Publishers Inc. ISBN 978-1-55860-032-4. URL <https://dl.acm.org/doi/10.5555/112922.112954>.
- Stuart Russell and Peter Norvig. *Artificial Intelligence - A Modern Approach*. Prentice Hall Series. Pearson Education, third edition, 2010.
- Markus Schlosser. Agency. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2019 edition, 2019. URL <https://plato.stanford.edu/archives/win2019/entries/agency/>.

- Steven Sloman, Aron K. Barbey, and Jared M. Hotelling. A Causal Model Theory of the Meaning of Cause, Enable, and Prevent. *Cognitive Science*, 33(1):21–50, 2009. ISSN 1551-6709. doi: 10.1111/j.1551-6709.2008.01002.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1551-6709.2008.01002.x>. \_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1551-6709.2008.01002.x>.
- Suzanne Tolmeijer, Markus Kneer, Cristina Sarasua, Markus Christen, and Abraham Bernstein. Implementations in Machine Ethics: A Survey. *ACM Computing Surveys*, 53(6):132:1–132:38, December 2021. ISSN 0360-0300. doi: 10.1145/3419633. URL <https://dl.acm.org/doi/10.1145/3419633>.
- Richard W. Wright. Causation in Tort Law. *California Law Review*, 73(6):1735–1828, 1985. ISSN 0008-1221. doi: 10.2307/3480373. URL <https://www.jstor.org/stable/3480373>. Publisher: California Law Review, Inc.
- Richard W. Wright. Causation, Responsibility, Risk, Probability, Naked Statistics, and Proof: Pruning the Bramble Bush by Clarifying the Concepts. *Iowa Law Review*, 73:1001, December 1988. URL [https://scholarship.kentlaw.iit.edu/fac\\_schol/698](https://scholarship.kentlaw.iit.edu/fac_schol/698).
- Richard W. Wright. The NESS Account of Natural Causation: A Response to Criticisms. In R Goldberg, editor, *Perspectives on Causation*. Social Science Research Network, Rochester, NY, hart publishing edition, July 2011. URL <https://papers.ssrn.com/abstract=1918405>.