

A Formalization of Kant’s Second Formulation of the Categorical Imperative

Martin Mose Bentzen and **Felix Lindner**

Management Engineering, Technical University of Denmark, Lyngby, Denmark
Department of Computer Science, University of Freiburg, Freiburg im Breisgau, Germany

Abstract

We present a formalization and computational implementation of the second formulation of Kant’s categorical imperative. This ethical principle requires an agent to never treat someone merely as a means but always also as an end. Here we interpret this principle in terms of how persons are causally affected by actions. Introducing Kantian causal agency models in which moral patients, actions, goals, and causal influence are represented. We show how to formalize several readings of Kant’s categorical imperative that correspond to Kant’s concept of strict and wide duties towards oneself and others. Stricter versions handle cases where an action directly causally affects oneself or others, whereas the wide version maximizes the number of persons being treated as an end. We discuss limitations of our formalization by pointing to one of Kant’s cases that the machinery cannot handle in a satisfying way.

Introduction

It has been suggested that artificial agents, such as social robots and software bots, must be programmed in an ethical way in order to remain beneficial to human beings. One prominent ethical theory was proposed by Immanuel Kant (1785). Here, we propose a formalization and implementation of Kant’s ethics with the purpose of guiding artificial agents that are to function ethically. In particular, the system will be able to judge whether actions are ethically permissible according to Kant’s ethics. In order to accomplish this we focus on the second formulation of Kant’s categorical imperative. Kant proposed three formulations of the categorical imperative. We formalize and implement the second formulation and do not take a stance on the interrelation of Kant’s three formulations.

The second formulation of Kant’s categorical imperative reads:

Act in such a way that you treat humanity, whether in your own person or in the person of any other, never merely as a means to an end, but always at the same time as an end. (Kant, 1785)

We take it to be the core of the second formulation of the categorical imperative that all rational beings affected by our actions must be considered as part of the goal of the action.

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The paper is structured as follows: We first briefly review related work. Then, building upon our earlier work (Lindner, Bentzen, and Nebel 2017), we introduce an extension of Pearl-Halpern-style causal networks which we call Kantian causal agency models. These models serve as a formal apparatus to model the morally relevant aspects of situations. We then define an action’s permissibility due to the categorical imperative, while considering two readings of *being treated as a means*. To deal with Kant’s wider duties, we introduce an extra condition according to which an agent should maximize the number of persons being treated as an end. Finally, we briefly showcase the computational implementation of the categorical imperative within the HERA software library¹.

Related work

In machine ethics, several ethical theories have been formalized and implemented, e.g., utilitarianism, see (Horty 2001; Arkoudas, Bringsjord, and Bello 2005), the principle of double effect, see (Bentzen 2016; Govindarajuli and Bringsjord 2017), pareto permissibility, see (Lindner, Bentzen, and Nebel 2017), and Asimov’s laws of robotics, see (Winfield, Blum, and Liu 2014).

It has been suggested for some time that Kant’s ethics could be formalized and implemented computationally, see (Powers 2006; Abney 2012). Powers (2006) suggests three possible ways of formalizing Kant’s first formulation of the categorical imperative, through deontic logic, non-monotonic logic, or belief revision. The first formulation of the categorical imperative states that you must be able to want that the subjective reasoning (or maxim) motivating your action becomes a universal law and as Kant claims that this in some cases is a purely formal matter, it should be possible to formalize it. However, Powers does not provide details of a formalization or a computational implementation, so the formalization of the first formulation in effect remains an open problem.

The work presented here differs in that we focus on the second formulation of the categorical imperative and in that we present a precise formal representation and computational implementation of the formal theory. Rather than taking a starting point in one of the paradigms Powers suggests,

¹<http://www.hera-project.com>

we use formal semantics and causal agency modelling as this is fitting for the means-end reasoning central to the second formulation. Philosophically, our formalization is best seen as a rational reconstruction within this framework of what we take to be the central ideas of Kant’s second formulation.

Ultimately, although we are sympathetic to Kant and sensitive to the original text, the goal of our work is not to get close to a correct interpretation of Kant, but to show that our interpretation of Kant’s ideas can contribute to the development of machine ethics. To meet this goal, our interpretation has to be detailed and explicit enough to provide a decision mechanism for the permissibility or not of specific actions in specific situations.

Kantian Causal Agency Models

In order to formalize the second formulation of the categorical imperative, we assume some background theory. First, we assume that *actions* are performed by *agents*, and that actions and their *consequences* can *affect* a set of *moral patients*, i.e. persons who must be considered ethically in a situation. The agent itself is also one of the moral patients. The agent has available a set of actions which will have consequences given *background conditions*. Some of the action’s consequences are the *goals* of the action. The actions and consequences that together cause my goal are the *means* of the action. Patients, who are affected by these means are *treated as a means*, and patients, who are affected by my goal are *treated as an end*. For example, I (agent) have the option available to press the light switch (action), and given that the light bulb is not broken (background condition), the light will go on (consequence), which leads to me being able to read my book (consequence). The last consequence was also my goal, and it affects me in a positive way. The action thus treats me as an end.

Within this informally characterized framework, we can reformulate the second formulation of the categorical imperative as follows:

Act in such a way, that whoever is treated as a means through your action (positively or negatively and including yourself), must also be treated as an end of your action.

The purpose of what follows is to formalize these intuitions. As a first step, we now give the formal definition of the models we will be using in Def. 1. We call these models *Kantian causal agency models* to set them apart from the causal agency models we used in our earlier work (Lindner, Bentzen, and Nebel 2017), and which had no formal tools to consider moral patients affected by one’s actions.

Definition 1 (Kantian Causal Agency Model)

A Kantian causal agency model M is a tuple (A, B, C, F, G, P, K, W) , where A is the set of action variables, B is a set of background variables, C is a set of consequence variables, F is a set of modifiable boolean structural equations, $G = (Goal_1, \dots, Goal_n)$ is a list of sets of variables (one for each action), P is a set of moral patients (includes a name for the agent itself), K is the ternary affect relation $K \subseteq (A \cup C) \times P \times \{+, -\}$, and

W is a set of interpretations (i.e., truth assignments) over $A \cup B$.

A (actions), B (background variables) and C (consequences) are finite sets of boolean variables with B and C possibly empty. W is a set of boolean interpretations of $A \cup B$. Thus, the elements of W set the truth values of those variables that are determined externally, and thus specify the concrete situation. We require that all interpretations in W assign true to exactly one action $a \in A$. As a notational convention, by M, w_a and M, w_b we distinguish two situations that only differ in that in the first situation, action a is performed, and in the second, action b is performed.

Causal influence is determined by the set F of boolean-valued structural equations. Each variable $c_i \in C$ is associated with the function $f_i \in F$. This function will give c_i its value under an interpretation $w \in W$. An interpretation w is extended to the consequence variables as follows: For a variable $c_i \in C$, let $\{c_{i1}, \dots, c_{im-1}\}$ be the variables of $C \setminus \{c_i\}$, $B = b_1, \dots, b_k$, and $A = \{a_1, \dots, a_n\}$ the action variables. The assignment of truth values to consequences is determined by:

$$w(c_i) = f_i(w(a_1), \dots, w(a_n), w(b_1), \dots, w(b_k), w(c_{i1}), \dots, w(c_{im-1}))$$

To improve readability, we will use the notation $c := \phi$ to express that c is true if ϕ is true, where ϕ can be any boolean formula containing variables from $A \cup B \cup C$ and its negations. For instance, the boolean structural equations for the light-switch example will be written as $F = \{lightOn := press \wedge \neg bulbBroken, canReadBook := lightOn\}$.

In the general setting, it may be unfeasible to extend an interpretation from the action variables to the rest of the variables, because it is possible that the value of some variable depends on the value of another variable, and the value of the latter variable depends on the value of the former. Dependence is defined in Def. 2.

Definition 2 (Dependence)

Let $v_i \in C, v_j \in A \cup B \cup C$ be distinct variables. The variable v_i depends on variable v_j , if, for some vector of boolean values, $f_i(\dots, v_j = 0, \dots) \neq f_i(\dots, v_j = 1, \dots)$.

Following Halpern (2016), we restrict causal agency models to acyclic models, i.e., models in which no two variables are mutually dependent on each other. First, note that the values of action variables in set A are determined externally by the interpretations in W . Thus, the truth values of action variables do not depend on any other variables. Additionally, we require that the transitive closure, \prec , of the dependence relation is a partial order on the set of variables: $v_1 \prec v_2$ reads “ v_1 is causally modified by v_2 ”. This enforces absence of cycles. In case of acyclic models, the values of all consequence variables can be determined unambiguously: First, there will be consequence variables only causally modified by action variables, and whose truth value can thus be determined by the values set by the interpretation. Call these consequence variables *level one*. On *level two*, there will be consequence variables causally modified by actions and

level-one consequence variables, and so on (Bentzen 2016; Halpern 2016).

Some of the definitions below will make use of causality. Thus, to take causation into account, definition 3 defines the relation of Y being a but-for cause of ϕ , see (Halpern 2016). Definition 3 makes use of *external interventions* on models. An external interventions X consists of a set of literals (viz., action variables, consequence variables, background variables, and negations thereof). Applying an external intervention to a causal agency model results in a new causal agency model M_X . The truth of a variable $v \in A \cup B \cup C$ in M_X is determined in the following way: If $v \in X$, then v is true in M_X , if $\neg v \in X$, then v is false in M_X , and if neither $v \in X$ nor $\neg v \in X$, then v is true in M_X if and only if v is true in M . External interventions thus override structural equations of the variables occurring in X .

Definition 3 (Actual But-For Cause)

Let y be a literal and ϕ a formula. We say that y is an actual but-for cause of ϕ (notation: $y \rightsquigarrow \phi$) in the situation the agent choses option w_a in model M , if and only if $M, w_a \models y \wedge \phi$ and $M_{\{\neg y\}}, w_a \models \neg \phi$.

The first condition says that both the cause and the effect must be actual. The second condition says that if y had not been the case, then ϕ would have not occurred. Thus, in the chosen situation, y was necessary to bring about ϕ . Sometimes it is convenient to be able to talk about causality in the logical language used to talk about the models. In accordance with Halpern’s (2016) notation, we can define that y is a but-for cause of ϕ in model M, w_a iff $M, w_a \models y \wedge \phi \wedge [y \leftarrow \perp] \neg \phi$ holds. This definition is equivalent to the definition in Def. 3. So, in the above light-switch example, we might want to determine if $\neg \text{bulbBroken}$ is a cause of canReadBook . And indeed, if we consider the model $M_{\text{bulbBroken}}, w_{\text{press}}$, we can determine that canReadBook is false, written $M_{\text{bulbBroken}}, w_{\text{press}} \not\models \text{canReadBook}$. Likewise, we obtain $M, w_{\text{press}} \models [\text{bulbBroken} \leftarrow \top] \neg \text{canReadBook}$, i.e., under the intervention that sets bulbBroken true, canReadBook is false.

The concept of but-for cause allows us to introduce the useful concept of *direct consequences* (Def. 4).

Definition 4 (Direct Consequence)

A consequence $c \in C$ is a direct consequence of $v \in A \cup B \cup C$ in the situation M, w_a iff $M, w_a \models v \rightsquigarrow c$.

Persons can be affected by actions or consequences either in a positive or in a negative way. To represent that some action or consequence affects a person positively or negative, we introduce the notations \triangleright_+ and \triangleright_- , respectively. Thus, $M, w_a \models c \triangleright_+ p$ holds iff $(a, c, +) \in K$, and $M, w_a \models c \triangleright_- p$ holds iff $(a, c, -) \in K$. We use \triangleright in case the valence of affection is not relevant.

This finalizes the exposition of the background theory. We will now consider how to make permissibility judgments about actions as defined in the context of Kantian causal agency models using the categorical imperative.

Categorical Imperative Defined

The second formulation of the categorical imperative requires an agent to never treat someone merely as a means

but always also as an end. Thus, to formalize under which conditions an action is permitted by the categorical imperative, we first define the concept of someone being *treated as an end* (Def. 5), and then formalize two possible readings of the concept of someone being *treated as a means* (Def. 6 and Def. 7).

Definition 5 (Treated as an End)

A patient $p \in P$ is treated as an end by action a , written $M, w_a \models \text{End}(p)$, iff, the following conditions hold:

1. Some goal g of a affects p positively
 $M, w_a \models \bigvee_g (G(g) \wedge g \triangleright_+ p)$.
2. None of the goals of a affect p negatively
 $M, w_a \models \bigwedge_g (G(g) \rightarrow \neg(g \triangleright_- p))$

Thus, being treated as an end by some action means that some goal of the action affects one in a positive way. One could say that the agent of the action, by performing that action, considers those who benefit from his goal. Things are less clear regarding the concept ‘being treated as a means’. As a first step, we define two versions of the concept which we refer to as *Reading 1* and *Reading 2*. Both readings make use of the causal consequences of actions. Reading 1 considers a person used as a means in case she is affected by some event that causally brings about some goal of the action. As a consequence, negative side effects are permitted. Consider, for instance, the classical trolley dilemma, where the agent has the choice to either pull the lever to lead the tram onto the second track killing one person, or refraining from pulling letting the tram kill five persons on the first track. Under Reading 1, in case of pulling, the one agent is not treated as a means. This reading is probably closest to what we informally mean by ‘being treated as a means’. Reading 2 requires that everybody affected by any direct consequence of the action is considered as a goal, hence, also the person on the second track. As a consequence, everyone treated as a means according to Reading 1 is also treated as a means according Reading 2, but Reading 2 may include additional patients. This reading is further from the everyday understanding of means-end reasoning, but is probably closer to what some people expect of a Kantian ethics. We consider it a feature of a formal framework that we are able to make such clear and meaningful distinctions, but we leave it for the modeler to decide which one of the readings is more useful for a given application.

Definition 6 (Treated as a Means, Reading 1)

A patient $p \in P$ is treated as a means by action a (according to Reading 1), written $M, w_a \models \text{Means}_1(p)$, iff there is some $v \in A \cup C$, such that v affects p , and v is a cause of some goal g , i.e., $M, w_a \models \bigvee_v ((a \rightsquigarrow v \wedge v \triangleright p) \wedge \bigvee_g (v \rightsquigarrow g \wedge G(g)))$.

Definition 7 (Treated as a Means, Reading 2)

A patient $p \in P$ is treated as a means by action a (according to Reading 2), written $M, w_a \models \text{Means}_2(p)$, iff there is some direct consequence $v \in A \cup C$ of a , such that v affects p , i.e., $M, w_a \models \bigvee_v (a \rightsquigarrow v \wedge v \triangleright p)$.

Having defined both *being treated as an end* and *being treated as a means*, the permissibility of actions according to the second formulation of the categorical imperative

can now be defined in Def. 8. The formulation requires that noone is merely used as a means, but always at the same time as an end.

Definition 8 (Categorical Imperative)

An action a is permitted according to the categorical imperative, iff for any $p \in P$, if p is treated as a means (according to Reading N) then p is treated as an end
 $M, w_a \models \bigwedge_{p \in P} (Means_N(p) \rightarrow End(p))$

There are thus two main reasons why an action is not permitted. Either a patient is treated as a means but is left out of consideration by the end of the action. Or, the the action is done for an end that affects someone negatively.

Note that the choice between Reading 1 and Reading 2 yields significant implications, as can be seen in the Trolley dilemma: According to Reading 1, pulling the lever is permissible, because the one person on the other track is not treated as a means, and therefore it is permissible to not consider her as an end. According to Reading 2, however, pulling the lever is impermissible, because the one person on the other track is treated as a means while she does not benefit from the goal (and is thus not treated as an end). Refraining from action is permitted by the categorical imperative according to both the readings of ‘being treated as a means’ (more about this issue below in the Section on wide duties).

Cases of Strict Duty

We will now provide examples that highlight aspects of the definition of the categorical imperative. Although these do not prove it correct in any formal sense they can be used to discuss its appeal as an ethical principle as an explication of Kant’s ideas. First, we rephrase three cases that contain what Kant calls strict duties (and two of which Kant himself used to explain his ideas).

Example 1: Suicide

Bob wants to commit suicide, because he feels so much pain he wants to be relieved from. This case can be modeled by a causal agency model M_1 that contains one action variable *suicide* and a consequence variable *dead*. Death is the goal of the suicide action (as modeled by G), and the suicide affects Bob (as modeled by K). In this case, it does not make a difference whether the suicide action affects Bob positively or negatively.

$$\begin{aligned} A &= \{suicide\} \\ C &= \{dead\} \\ F &= \{dead := suicide\} \\ K &= \{(suicide, Bob, +)\} \\ G &= (Goal_{suicide} = \{dead\}) \end{aligned}$$

We assume that the suicide affects noone other than Bob, because Kant’s argument is not about the effect of suicide on other people but about the lack of respect of the person committing suicide. The reason why Bob’s suicide is not permitted is that the person affected by the suicide, viz., Bob, does not benefit from the goal, because he is destroyed and thus cannot be affected positively by it. He is thus treated as

a means to his own annihilation from which he receives no advancement. Therefore, the first condition of the categorical imperative (Def. 8) is violated according to both readings (1 and 2), because $M_1, w_{suicide} \models Means_{\{1,2\}}(Bob)$ holds but $M_1, w_{suicide} \not\models End(Bob)$ does not.

As noted above, we could say that the suicide affects Bob negatively, and the action would also be impermissible. The reason for the impermissibility of suicide also in this case is not due to the fact that Bob does something harmful towards himself. As Kant also remarks, other harmful actions would be allowed, e.g., risking your life or amputating a leg to survive. To see this, M_1 can be slightly modified to M_1^* : We rename *suicide* to *amputate* and *dead* to *survives*. Moreover, $K = \{(amputate, Bob, -), (survives, Bob, +)\}$, $G = (Goal_{amputate} = \{survives\})$. In this case, Bob is positively affected by the goal, and thus the act of amputation is permitted. The modified example also shows that in some cases, the categorical imperative is more permissive than the principle of double effect, which strictly speaking never allows negative means to an end.

Example 2: Giving Flowers

We consider the fact that an action can be judged as impermissible by the categorical imperative although noone is negatively affected a feature of the categorical imperative that inheres in no other moral principles formalized so far. The following example showcases another situation to highlight this feature: Bob gives Alice flowers in order to make Celia happy when she sees that Alice is thrilled about the flowers. Alice being happy is not part of the goal of Bob’s action. We model this case by considering a Kantian causal agency model M_2 :

$$\begin{aligned} A &= \{give_flowers\} \\ C &= \{alice_happy, celia_happy\} \\ P &= \{Bob, Alice, Celia\} \\ F &= \{alice_happy := give_flowers \\ &\quad celia_happy := alice_happy\} \\ K &= \{(alice_happy, Alice, +), \\ &\quad (celia_happy, Celia, +)\} \\ G &= (Goal_{give_flowers} = \{celia_happy\}) \end{aligned}$$

The action *give_flowers* is not permitted according to the categorical imperative, because Bob is using Alice as a means to make Celia happy, but not considering her as part of the goal of the action. This action is immoral, even though the action has positive consequences for all, and no bad consequence are used to obtain a good one and thus this example shows how the Kantian principle differs from other ethical principles such as utilitarianism and the principle of double effect, because these principles would permit the action. The model M_2 can easily be turned into model M_2^* , in which Bob’s action is permitted by the Kantian principle. The only thing in which M_2^* differs from M_2 is that the variable *alice_happy* is added to the set $Goal_{give_flowers}$. In this case, Alice is both treated as a means and treated as an end, which is permitted by the categorical imperative.

This example demonstrates how demanding the categorical imperative is, because the principles requires that everybody affected by ones' action must be treated as a goal: This includes the taxi driver that drives you to your destination, as well as the potential murderer you defend yourself against. In these examples, one is required to, e.g., have the taxi driver's earning money among one's goals, and the murderer's not going to jail.

Example 3: False Promise

We return to a case mentioned by Kant himself. Consider that Bob makes a false promise to Alice. Bob borrows one 100 Dollars from Alice with the goal of keeping the money forever. He knows that it is an inevitable consequence of borrowing the money that he will never pay it back. Hence, the model of this situation, M_3 , contains the following items:

$$\begin{aligned}
 A &= \{borrow\} \\
 C &= \{bob_keeps_100Dollar_forever\} \\
 P &= \{Alice, Bob\} \\
 F &= \{bob_keeps_100Dollar_forever := borrow\} \\
 K &= \{(borrow, Bob, +), (borrow, Alice, -), \\
 &\quad (bob_keeps_100Dollar_forever, Bob, +), \\
 &\quad (bob_keeps_100Dollar_forever, Alice, -)\} \\
 G &= (Goal_{borrow} = \{bob_keeps_100Dollar_forever\})
 \end{aligned}$$

The action is impermissible, because Alice is treated as a means (by both Reading 1, Def. 6, and Reading 2, Def. 7) while she is not treated as an end. In this case, none of the two conditions for 'being treated as an end' (Def. 5) are met.

Cases of Wide Duty

Examples 1, 2 and 3 are instances of what Kant calls necessary, strict, narrower duties to oneself and to others, and it seems obvious they involve using a person as a means. Kant also presents two other examples to which we now turn. These involve what Kant calls contingent, meritorious, or wider duties. His arguments for these appear more vague and at least from our perspective harder to handle. We now turn to wide duties and discuss, through an example, how actions that indirectly affect others by refraining from preventing harmful consequences could be handled, and we learn through another example where the limitations of our formalization attempt are.

Example 4: Not Helping Others

Bob who has everything he needs, does not want to help Alice who is in need. Let us assume she is drowning and Bob is refraining from saving her live. Formally, the situation in the example can be represented with a causal agency model M_4 that contains one background variable *accident* representing the circumstances that led to Alice being in dire straits, two action variables *rescue* and *refrain* and a consequence

variable *drown*. Moreover, \neg *drown* is the goal of *rescue*.

$$\begin{aligned}
 A &= \{rescue, refrain\} \\
 B &= \{accident\} \\
 C &= \{drown\} \\
 P &= \{Alice, Bob\} \\
 F &= \{drown := accident \wedge \neg rescue\} \\
 K &= \{(drown, Alice, -), (\neg drown, Alice, +)\} \\
 G &= (Goal_{rescue} = \{\neg drown\}, Goal_{refrain} = \emptyset)
 \end{aligned}$$

According to the categorical imperative using Readings 1 and 2 of 'being treated as a means' both *rescue* and *refrain* are permitted. Bob is strictly speaking not using Alice as a means by going about his business. Kant gives us a clue of how to formalize an argument against refraining in that he says we have to make other people's ends our own as far as possible. Kant writes that 'For a positive harmony with humanity as an end in itself, what is required is that everyone positively tries to further the ends of others as far as he can.' One way of understanding this is as an additional requirement on top of the categorical imperative of choosing an action whose goals affect most people positively.

Definition 9 (Meritorious principle)

Among actions permitted by the categorical imperative, choose one whose goals affect most patients positively.

The meritorious principle thus goes beyond simply avoiding to treat others as means by acitvely helping them. As formulated here, the principle is compatible with the categorical imperative. In example 4, it requires of the agent to choose saving Alice, because the goal advances her. There may be several actions advancing the same number of agents, in which case the agent can choose freely (or randomly) amongst them. One could also take Kant to imply a second condition to the meritorious principle, to prevent as many people being negatively affected by circumstances as possible. In the current example, both conditions would lead to the same result.

Unhandled Case: Not Using Your Talent

As a final example, consider the following situation: Bob has the talent to become a great artist. However, he wonders whether it is permissible to just be lazy and enjoy life instead of working hard to improve himself. Strictly speaking Bob is not working to anyone's disadvantage by being lazy and thus the definitions of 'being treated as a means' advanced above will not cover this example. As the goal of enjoying life and the goal of making art both benefit Bob, the meritorious principle also cannot be used to make the distinction. What Kant says is that laziness could be consistent with the preservation of humanity but does not harmonize with its advancement. He also writes that a rational being necessarily wills that all his capacities are developed. However, it is not clear to us what consistutes the advancement of humanity beyond the sheer feeling of happiness. The example is further complicated by the fact that Kant says that this is a duty one has towards oneself, not others. Therefore, it would be inappropriate to solve this case by introducing others into the model that would benefit from Bob becoming an artist.

In the current formalization, we have no means to represent the relevant aspects that render laziness impermissible and becoming an artist permissible for the right reasons. We thus take this example to showcase a limitation of our treatment of Kant’s ethics, and leave a formalization that could capture this last example for further research.

Implementation

The formalization of Kant’s second formulation of the categorical imperative has been implemented within the Hybrid Ethical Reasoning Agent software library (short: HERA).² The general goal of the HERA project is to provide theoretically well-founded and practically usable logic-based machine ethics tools for implementation in artificial agents, such as companion robots with moral competence (Lindner and Bentzen 2017; Lindner, Wächter, and Bentzen 2017). The core of HERA consists of a model checker for causal agency models. Thus, the ethical principles are implemented as sets of logical formulas, just the same way we use logical formulas to define the conceptual notions. To showcase the use the categorical imperative from a Python program, Listing 1 reconsiders a representation of the suicide case.

```
{
  "actions": ["suicide"],
  "background": [],
  "consequences": ["dead"],
  "patients": ["Bob"],
  "mechanisms": {"dead": "suicide"},
  "affects": {"suicide": [["Bob", "+" ]],
              "dead": []},
  "goals": {"suicide": ["dead"]}
}
```

Listing 1: A sample JSON encoding of the suicide case.

The workflow for using HERA requires to first generate a causal agency model like the one in Listing 1. Given such a model, arbitrary logical formulae can be checked for being satisfied or not by this model. This way, the conditions of ethical principles like the Kantian categorical imperative as defined in Def. 8 can be checked for satisfaction.

To support the usage of the HERA library, the logical formulae to be checked for ethical principles already included in HERA are encapsulated into prepared classes. Listing 2 shows a sample interaction. The first three commands load the implementations of two syntactical entities of the logical language (the predicates `Means` and `End`), the causal agency model from the `semantics` package, and the categorical imperative using Reading 1 of ‘being treated as a means’ from the `principles` package. The third command loads the suicide example and sets the external variable `suicide True`. This way, the `suicide` action is chosen in the situation, and the truth values of the consequence variables can be evaluated the way explained in the section on Kantian causal agency models. In the concrete case, `True` will be assigned to the variable `dead`. The fourth command

²The HERA software is available from <http://www.hera-project.com>. It is fully implemented in Python and can be installed via the PyPI repository (package name: `ethics`).

asks whether, in the resulting situation, Bob is used as a means according to Reading 1 (see Def. 6). The answer is `True`, because Bob is affected by the action (`suicide`) and the action is a but-for cause of Bob’s goal (`dead`). The fifth command asks if Bob is used as an end. This query returns `False`, because Bob is not affected by the goal (see Section Example 1: Suicide). All in all, the action is not permissible according to the categorical imperative, and the output of command six is accordingly.

```
from ethics.language import Means, End
from ethics.semantics import
    CausalModel as cm
from ethics.principles import
    KantianHumanityPrincipleReading1 as ci
m = cm("suicide.json", {"suicide": True})
m.models(Means("Reading-1", "Bob"))
output: True
m.models(End("Bob"))
output: False
m.evaluate(ci)
output: False
```

Listing 2: A sample interaction with the Python package `ethics`, which we develop and maintain as the standard implementation of HERA.

Conclusion

We have shown proof of principle that Kant’s second formulation of the categorical imperative can be formalized and implemented computationally. The strict duties towards yourself and others are defined, given goals, causal mechanisms, and the affects relation. To define permissibility according the categorical imperative, we have defined ‘being treated as an end’, and we formalized two readings of ‘being treated as a means’ that meet different intuitions about this concept. The formalization deals well with Kant’s own examples of strict duties. We were also able to partly deal with Kant’s wide duties by defining an additional condition that requires agents to maximize the number of persons being treated as an end.

We envision that the theory will be used as a tool for the comparison of morally relevant aspects of different views on morally delicate cases, thus helping people to have moral discussions. Moreover, we aim at allowing automatic moral judgments in line with Kant in robots such as self-driving cars and care robots. Our future research will investigate whether and under which circumstances Kantian reasoning the way it is presented here is perceived as appropriate for social robots as compared to other types of moral reasoning.

References

- Abney, K. 2012. Robotics, ethical theory, and metaethics: A guide for the perplexed. In Lina, P.; Abney, K.; and Bekey, G. A., eds., *Robot Ethics: The Ethical and Social Implications of Robotics*. MIT Press. 35–52.
- Arkoudas, K.; Bringsjord, S.; and Bello, P. 2005. Toward ethical robots via mechanized deontic logic. Technical report, AAAI Fall Symposium on Machine Ethics, AAAI.

- Bentzen, M. 2016. The principle of double effect applied to ethical dilemmas of social robots. In *Robophilosophy 2016/TRANSOR 2016: What Social Robots Can and Should Do*. IOS Press. 268–279.
- Govindarajuli, N. S., and Bringsjord, S. 2017. On automating the doctrine of double effect. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*. 4722–4730.
- Halpern, J. Y. 2016. *Actual Causality*. The MIT press.
- Horty, J. F. 2001. *Agency and Deontic Logic*. Oxford University Press.
- Kant, I. 1785. *Grundlegung zur Metaphysik der Sitten*. Felix Meiner Verlag, seventh edition.
- Lindner, F., and Bentzen, M. 2017. The hybrid ethical reasoning agent IMMANUEL. In *Proceedings of the Companion 2017 Conference on Human-Robot Interaction (HRI)*. IEEE. 187–188.
- Lindner, F.; Bentzen, M.; and Nebel, B. 2017. The HERA approach to morally competent robots. In *Proceedings of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE/RSJ.
- Lindner, F.; Wächter, L.; and Bentzen, M. 2017. Discussions about lying with an ethical reasoning robot. In *Proceedings of the 2017 IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE.
- Powers, T. M. 2006. Prospects for a kantian machine. *IEEE Intelligent Systems* 21(4):46–51.
- Winfield, A. F.; Blum, C.; and Liu, W. 2014. Towards an ethical robot: internal models, consequences and ethical action selection. In Mistry, M.; Leonardis, A.; M. Witkowski; and Melhuish, C., eds., *Advances in Autonomous Robotics Systems*. Springer. 85–96.