

Taking Principles Seriously: A Hybrid Approach to Value Alignment

Tae Wan Kim¹, John Hooker¹, and Thomas Donaldson²

¹Carnegie Mellon University, USA

²University of Pennsylvania, USA

December 20, 2020

Abstract

An important step in the development of value alignment (VA) systems in AI is understanding how VA can reflect valid ethical principles. We propose that designers of VA systems incorporate ethics by utilizing a hybrid approach in which both ethical reasoning and empirical observation play a role. This, we argue, avoids committing “naturalistic fallacy,” which is an attempt to derive “ought” from “is,” and it provides a more adequate form of ethical reasoning when the fallacy is not committed. Using quantified model logic, we precisely formulate principles derived from deontological ethics and show how they imply particular “test propositions” for any given action plan in an AI rule base. The action plan is ethical only if the test proposition is empirically true, a judgment that is made on the basis of empirical VA. This permits empirical VA to integrate seamlessly with independently justified ethical principles.

1 Introduction

Artificial intelligence (AI) technologies increasingly replace human decision makers. Worries rise about the compatibility of AI and human values. A growing number of researchers are examining how AI can acquire moral intelligence (Wallach and Allen, 2008; Burton et al., 2017; Walsh et al., 2019; Lin et al., 2011; Bringsjord, 2013; Scheutz and Arnold, 2016; Arnold et al., 2017; Arnold and Scheutz, 2018). Such attempts are recently lumped under the term “value alignment” (hereafter VA). Russell et al. (2015) highlight the need for VA and identify two options for achieving it:

“[A]ligning the values of powerful AI systems with our own values and preferences ... [could involve either] a system [that] infers the preferences of another rational or nearly rational actor by observing its behavior ... [or] could be explicitly inspired by the way humans acquire ethical values.”

As this passage suggests, one option for VA is to teach machines human preferences, and another is to teach machines ethics. The word “values” in fact has this double meaning. It can refer to what humans value in the sense of what they see as subjectively preferable, or it can refer to reasonably defensible ethical principles. The distinction is important, because we acquire knowledge of the two types of values in different ways.

A similar distinction occurs in previous literature under the names *top-down* and *bottom-up* VA (Allen et al., 2000; Allen, 2002; Allen et al., 2005, 2006; Wallach and Allen, 2008; Wallach et al., 2008). Russell et al. suggest a bottom-up approach in the form of inverse reinforcement learning, which allows a machine to internalize a pattern of preferences by observing how humans actually behave (Abbeel and Ng, 2004; Ng and Russell, 2000). Reinforcement learning, and machine learning (ML) in general, offer a number of advantages but must deal with such issues as inadequate reward functions to represent complex ethical norms, biased data, and opaqueness (Arnold et al., 2017; Prince and Pinker, 1988; Marcus, 2018). A promising alternative to ML is *logic-based* VA, which has received less attention despite having a long research record (Arkoudas et al., 2005; Bringsjord et al., 2006; Bringsjord and Taylor, 2012; Bringsjord, 2017; Govindarajulu and Bringsjord, 2017; Hooker and Kim, 2018).

In this paper, we make a case for *hybrid* VA that combines ML-based and logic-based approaches. A logic-based approach is especially important because it allows the use of “independently justified” or “independently defensible” ethical principles. By these we mean principles that find their justification in ethical theory. Such principles are “normative” in the sense commonly used by moral philosophers: they are prescriptive rather than descriptive and are elements of traditional normative moral theories such as deontology, consequentialism and virtue ethics. Such principles are increasingly discussed as candidates for computational use (Lindner and Bentzen, 2018; Lindner et al., 2020; Ganascia, 2007). Independently justified principles avoid many problems, including those associated with the well-known *is-ought gap*, one aspect of which is reflected in the unconscious biases now widely studied by behavioral ethicists (Bazerman and Tenbrunsel, 2011). In turn, we propose our own version of deontological VA for use in such a

Table 1: Comparison of ML-based and logic-based VA

	ML-based	Logic-based
Theory of mind	Connectionism	Computationalism
Base discipline	Statistics	Logic
AI techniques	Machine learning (automated statistics, deep learning)	Symbolic AI (i.e., GOFAI: Good Old Fashioned AI)
Value alignment	Bottom-up	Top-down
Example	ML system trained by lay people’s perception of fairness regarding autonomous vehicles and gender/racial discrimination.	Formalized normative principles (e.g., double effect theory, categorical imperatives) using symbolic logic (e.g., quantified modal logic)
Dual process theory	System 1	System 2

hybrid approach.

After elaborating on why a purely ML-based approach is inadequate, we show how symbolic logic enables the introduction of deontological reasoning into machine ethics. Rather than opting for a particular version of moral theory, we attempt to develop a comprehensive, ecumenical framework of ethical principles (Parfit, 2011). We first articulate univeralization, utilitarian, and autonomy-based principles in the idiom of quantified modal logic. We then use these principles to derive *test propositions*, also formulated in modal logic, for each action specified by an AI rule base. The action is ethical only if the test propositions are empirically true, a judgment that can be based on machine learning and empirical VA. This permits empirical VA to integrate seamlessly with independently justified ethical principles.

2 Two different value alignment systems

AI is an imitation game. It imitates the human mind. Because more than one theory of mind is possible, different models of AI are also possible, and so too, different models of VA. Broadly speaking, two categories of VA stand out, ML-based and logic-based, although neither is instantiated perfectly in any given working AI system (Table 1).

ML-based VA is connectionist. Connectionism holds that human intelligence can be explained and imitated by using artificial neural nets consisting of three kinds of connected units: input, hidden and output (Buckner and

Garson, 2019). Deep Learning (DL) exemplifies connectionism by utilizing a complex “automated statistics” based on a large number of hidden and opaque heuristics using associations (Danks, 2014). ML’s major advantage, which is especially obvious in an end-to-end model such as DL, is its powerful ability to imitate and further strengthen skill sets in training data. DL has illustrated the power of connectionist models by capably learning human skills, especially in the domain of pattern recognition, face recognition, medical diagnostic systems, and text reading.

2.1 The is-ought gap and the problem of bounded ethicality

Since connectionist systems are inductive, the quality of ML-based VA relies heavily on that of inputs. If training data is biased or unethical, the system will generate well-imitated, undesirable outputs. Microsoft’s AI-based chatterbot Tay (an acronym for “thinking about you”) was designed to engage with people on Twitter and learn from them how to carry on a conversation. When some people started tweeting racist and misogynistic expressions, Tay responded in kind. Microsoft immediately terminated the experiment (Wolf et al., 2017). Most algorithmic bias problems we see now are results of ML-based VA, which uses data sets from humans who already have implicit or explicit biases.

These mistakes reflect an error well-known to moral philosophers, the problem of deriving an “ought” from an “is,” sometimes called the “naturalistic fallacy.” From the fact that people behave in racist ways, it cannot follow that people ought to behave in such ways. While not a formal fallacy, the violation of the is-ought gap signals a form of epistemic naïveté, one that ignores the axiom in normative ethics that “no justifiable ‘ought’ can be derived directly from an ‘is.’” Disagreements about the robustness of the fallacy abound (Donaldson, 1994, 2012; Pigden, 2016; Woods and Maguire, 2017), and so this paper adopts a modest, workable interpretation of the is-ought gap coined recently by Daniel Singer, namely, “There are no valid arguments from non-normative premises to a relevantly normative conclusion” (Singer, 2015). Descriptive (or naturalistic) statements are reportive of what states of affairs are like, whereas normative statements are stipulative and action-guiding. Examples of the former are “The grass is green” and “Many people find deception to be unethical.” Examples of the latter are “You ought not murder” and “Lying is unethical.” Normative statements usually figure in the semantics of deontic (obligation-based) or evaluative expressions such as “ought,” “impermissible,” “wrong,” “good,” “bad,” or “unethical.” One may object that a high-level domain-general premise such as “machines ought

to have our values no matter what” might successfully link a descriptive premise to a normative conclusion. This objection is correct, but allows the original problem to pop up again at a deeper level. “What facts,” one might ask, “justify the conclusion that machines ought to imitate perfectly our behaviors?”

Using data from “unbiased” people’s behaviors seems an obvious solution, but the problem is more complicated than one thinks. The preceding decade of research in behavioral ethics has shown the existence of various pernicious influences on ethical decisions, often at an unconscious level. When these influences lead to unethical behavior that conflict with an actor’s moral beliefs and commitments (Moore et al., 2006), the phenomenon is often referred to as “bounded ethicality” (Bazerman, 2011; Bazerman and Tenbrunsel, 2011; Chugh et al., 2005; Tenbrunsel, 2005). One example of bounded ethicality is “ordinary prejudice,” which reveals itself in implicit associations about gender, race, and other demographic groups (Bertrand et al., 2005; Green et al., 2007; Greenwald et al., 2009; Rudman and Ashmore, 2007). These associations can lead to unintentionally discriminatory results, such as discriminatory hiring practices and unwarranted discrepancies in the evaluation of the skills and competencies of workers. Other elements of bounded ethicality include “in-group favoritism,” “self-serving bias,” and “motivated blindness,” the last of which refers to a systemic but unconscious failure to notice unethical behavior in oneself or others even when it is in one’s financial interest to do so (Bazerman and Moore, 2011; Moore et al., 2010). One might consider using professional moral philosophers’ opinions as training data for ML-based VA (Anderson and Anderson, 2011), but recent research shows both expert judgment generally and ethical expert judgment in particular to be frequently biased. Professional ethicists’ moral intuitions and specific judgements turn out to be as vulnerable to biases or irrelevant factors as those of lay persons (Schwitzgebel and Cushman, 2012; Wiegmann et al., 2020; Tobia et al., 2013; Schwitzgebel and Cushman, 2015; Egler and Ross, 2020). Because any attempt to use the ML-based VA system to generate the principles would be viciously circular, ML-based systems stand in need of independently defensible principles in order to evaluate even the training data to be used.

Logic-based VA is distinct from the ML-based in several ways. It is analogous to computationalism, in which human intelligence operates as a computer does, or in other words, in step with a set of systematic, abstract, symbol-and-rule mechanisms that are transparently expressed with formal-symbolic logic (Rescorla, 2020; Scheutz, 2002). Due to the popularity of ML systems, logic-based systems are sometimes referred to as GOFAI (“good old-fashioned AI”) (Haugeland, 1985). But logic-based AI is still widely

used, for instance, in the driving mechanisms of autonomous drones or cars, even though the pattern recognition mechanisms in these applications are primarily based on ML systems. Logic-based approaches are especially useful when formalizing independently defensible ethical principles of the sort invoked by professional philosophers. Such logic-based systems are sometimes labeled *symbolic AI*. Interestingly enough, formal logic is one of a few languages shared by both computer scientists and moral philosophers. Unlike eliminative (pure) connectionist systems, logic-based VA relies not on associations, but on deductive logic and logical proofs.

2.2 The problem of System 2 and systematicity

From a psychological perspective, ML systems are relevantly similar to what dual process theory (Kahneman, 2011) knows as “System 1” (Chauvet, 2018; Geffner, 2018; Rossi and Loreggia, 2019). It is opaque, fast, and intuitive to use. Dual process theory frames the human mind in terms of two distinctive processes: System 1 and System 2. In contrast to System 1 thinking, System 2 thinking is slow, transparent, analytical, logical, reasons-responsive and computational. Research shows that unethical and biased decisions are correlated with System 1 thinking, and that shifting the mode to System 2 thinking is often an effective way to avoid unethical behaviors (Bazerman and Gino, 2012; Bazerman and Sezer, 2016; Zhang et al., 2015; Sezer et al., 2015). This is despite the fact that System 1 thinking can be useful in other domains where intuitive associations are useful, such as in making heuristic decisions.

Because ML systems draw upon System 1 behavior, ML-based VA can be inherently vulnerable to unethical decision making. The “systematicity” challenge, neglected by connectionists for decades (Calvo and , eds.; Lake and Baroni, 2018; Alhama and Zuidema, 2019; Geffner, 2018; Marcus, 2001), sheds further light on this. In 1988, linguistic philosophers (Fodor and Pylyshyn, 1988) argued that connectionism confuses the intrinsically systematic nature of thought with a system of associations. More specifically, they argued that thoughts—e.g., “Mary loves John”—must involve operations with a set of rules (e.g., syntactic and semantic combinatorial relations or grammars). Pure or eliminative connectionist systems, which rely exclusively on associations, lack the ability to employ rules, and this seriously limits their ability to explain human thinking. A human who can think “Mary loves John” can also think “John loves Mary,” but purely connectionist systems trained by connectionist methods cannot systematically do the latter without further resources. Responding to this challenge, many connectionists have attempted

to show that structured ML systems might be redesigned, but the attempts underscore the eventual need for pure or eliminative ML systems that employ rule-like structures.

Our purpose here is not to adjudicate this debate. However, the debate itself reveals the need for connectionist systems to be used within their legitimate scope. In that sense, our view is roughly consistent with that of Paul Smolensky who responded to the systematicity challenge in his article, “On the proper treatment of connectionism (PTC)” (1988). Similar to dual process theory, Smolensky’s “proper treatment of connectionism” construes human intelligence in terms of two distinct realms: on the one hand, there is “cultural knowledge” (e.g., formalized knowledge presented by symbols and rule-like logic), and on the other, there is “individual knowledge” (e.g., perception, intuitive processing). Connectionist systems are adequate for the latter, but not the former. The proper treatment of connectionism entails that computational systems are necessary but insufficient for language-like processing because human language operates against a backdrop of empirical, common-sensical knowledge which, in turn, allows rules themselves to make sense.

This broad point is especially relevant for moral thought, in which the “reasoning” portion of moral thinking relies upon systemic operations instead of associations. A person who can reason, “It is wrong for Jane to gratuitously lie to Mary” can also reason “It is wrong for Mary to gratuitously lie to Jane” or “It is not wrong for . . .” Moral reasoning is fundamentally rule-based. It can be said that a person who concludes “It is wrong for Jane to lie to Mary” uses a rule such as “It is wrong for agent x to gratuitously lie to someone” and an empirical premise, “Jane gratuitously lies to Mary.”¹

3 Related work

Bringsjord and his collaborators (Bringsjord et al., 2006; Bringsjord and Taylor, 2012; Bringsjord, 2017; Arkoudas et al., 2005; Govindarajulu and Bringsjord, 2017) are the first we know to use deontic logic to explicitly represent philosophically justifiable ethical principles such as the doctrine of double effect. Our approach dovetails with that of Bringsjord in identifying the importance of deontic logic for teaching right and wrong to machines.

¹Moral particularism criticizes rule-based ethical theory, but grants easily that rules are used in moral reasoning, even as it critiques a one-size-fits-all approach. Interestingly, a rule-based or logic-based ethical theory is not committed to a rigorous one-size-fits-all approach (Smith and Dubbink, 2011).

Since his pioneering work, many others have attempted to represent ethical principles using deontic logic. These contributions reveal the versatility of deontic logic when formalizing not only deontological moral theory but other traditions, such as areteic theory (including virtue ethics) and commandment theory.

Our work is consistent with the established deontic tradition in moral philosophy that uses deontic logic to formalize deontological moral theory. Rather than opting for a particular version of moral theory, we attempt to develop a comprehensive, ecumenical framework of ethical principles (Parfit, 2011). We offer a deontological representation of three central ethical traditions, using a generalization principle, an autonomy principle, and a deontic utility principle. Using deontology, we indicate in outline how ethical obligations can be derived from first principles instead of relying on conflicting moral intuitions of what seems fair or unbiased. While ethical philosophy has been viewed as vague and subjective by the popular imagination, the deontological approach to moral philosophy is known for offering a rigorous foundation.

Allen et al. first suggested a hybrid approach to VA and recommended combining top-down and bottom-up approaches. Although their distinction can be more broadly construed, a typical top-down approach installs ethical principles directly into the machine, while a bottom-up approach typically asks the machine to learn prescriptive norms from experience. From an epistemological perspective, the typical bottom-up VA approach can result in teaching strategies that sometimes conflate “is” and “ought.” For example, one might suggest that a machine might learn ethics through a simulated process of evolution (Conitzer et al., 2017). The fact that certain ethical norms evolve does not imply that they are valid ethical principles (Berker, 2009; Nagel, 1979; McDowell, 1995; Rachels, 1990). It is true that bottom-up approach does not automatically commit the naturalistic fallacy, particularly if ethical principles validate the norms learned in this fashion (Wallach and Allen, 2008). Nonetheless, in our approach to hybrid VA, bottom-up learning does none of the normative work, but is used only to evaluate the truth of test propositions derived from ethical principles.

Another version of a hybrid approach to VA is advocated by Arnold et al., p. 81 who argue, “architectures must explicitly represent legal, ethical and moral principles,” while using them as principles for decision-making in order to achieve predictable decisions on the part of the system. Systems that uphold those principles as much as possible represent a more ethical path than systems that are less transparent less accountably trained, and less easily corrected. We largely agree with these authors, and our efforts

are indebted to their insightful criticism of the IRL-based VA. Arnold et al. suggest that the problems in the IRL approach can be significantly addressed by an hybrid approach in which explicitly written ethical rules can be imposed as constraints on what a machine learns from observation through IRL. We follow this very path by developing deontological principles as constraints, realizing nonetheless that one must ask precisely what remains within the unconstrained space of observational learning. If what remains is learning that includes ethical norms, then once again we confront the is-ought gap. If, on the other hand, it is learning that includes empirical facts about the world, then those facts alone cannot be transformed into “oughts.”

It is with this in mind that we offer a hybrid approach to VA that integrates independently justified ethical principles from the deontological tradition in ethics (Korsgaard, 1996; Nagel, 1986; O’Neill, 2014) with factual knowledge acquired through ML technology. Relevant facts may include observed preferences and values, but even such value-relevant facts cannot be the source of ethical principles.

Applying the imperative, “Thou shalt not kill,” to a given action requires at a minimum that someone knows the facts relevant to the action (Hare, 1991). The relevant facts, which may include observations of human values and preferences, do not by themselves decide what is ethical, but they factor into ethical assessment. In addition, action decisions almost always take the form, “If the facts are such-and-such, then do A,” which we refer to as an *action plan*. This provides a clue as to how VA can knit together empirical observation and ethical principles. The factual information in an action plan can be merged with ethical imperatives that depend on factual circumstances to arrive at an ethical judgment. The next section describes in detail how this can be accomplished.

4 Integrating ethical principles and empirical VA

We now show how deontologically derived ethical principles can combine with empirical facts in a systematic way. An adequate exposition of deontological reasoning is far beyond the scope of this paper, and we do not attempt to defend the specific ethical principles we have chosen, although we briefly explain why we think they are reasonable. Relevant literature is cited for readers who wish to study the underlying arguments in detail. Our purpose here is only to show how a careful statement of ethical principles clarifies how these principles can interrelate with empirical observation in VA.

We argue that expressing ethical assertions in the idiom of quantified

modal logic, as developed in Hooker and Kim (2018), makes the relationship between ethical principles and empirical observation perspicuous. Specifically: *ethical principles imply certain logical propositions that must be true in order for a given action plan to be ethical, and empirical observation determines whether these propositions are, in fact, true.* We refer these propositions as *test propositions*, whose empirical evaluation typically requires observation of human values, beliefs, and behavior. The test propositions need not appear alongside the action plans in an AI system, but they can be generated and evaluated automatically if desired (Section 4.5).

Thus the role of ethics in hybrid VA is to derive necessary conditions for the rightness of specific actions, and the role of empirical VA is to ascertain whether these conditions are satisfied in the real world.

4.1 Actions and Reasons

Deontology derives ethical principles from the logical structure of action (Kant, 1785; Wood, 1999; O’Neill, 2014; Hooker and Kim, 2019). It begins with the necessity of distinguishing free action from mere behavior, insofar as causally speaking, both are determined by chemical and physical forces. Contemporary deontological thinkers usually base the distinction between free and causally determined behavior on a Kantian *dual standpoint* theory of ethics that identifies free action as behavior for which the agent has *reasons* (Bilgrami, 1996; Korsgaard, 1996; Nagel, 1986; Nelkin, 2000). Such reasons are not themselves psychological causes or motivations, but considerations that the agent consciously makes to justify a choice. The reasons need not be good or convincing ones from another agent’s perspective, but must be sufficiently coherent to serve as an explanation of why the agent chose the action.

Ethical principles are necessary conditions for the coherence or intelligibility of the reasons behind an action. While a number of necessary conditions for coherence are possible, ethical principles rest on the *universality of reason*: an agent who takes a set of reasons as justifying an action must in order to be consistent take the reasons as justifying the same action for any agent to whom those reasons apply.

We focus on the three ethical principles that have been most intensely studied in the literature—generalization, utility maximization, and respect for autonomy. Each states a necessary condition for ethical conduct. We make no claim that they are exhaustive, but only that they illustrate how empirical VA can be anchored by ethical principles.

Before proceeding, two caveats are in order. First, in this paper we

do not attempt to convince readers of the superiority of the deontological tradition or its premise that principles can be discovered through an analysis of the logical structure of action. Our aim is more modest: to show that deontology is particularly suitable for hybrid VA. Two of the three principles we employ, generalization and respect for autonomy, have historical roots in Kant’s *The Formula of the Universal Law* and *The Formula of Humanity* (Wood, 1999), although our formulations of them differ. Second, we also use a deontic model of utilitarianism (e.g., Cummiskey, 1996) in order to make utilitarianism consistent with the other two other principles.

4.2 Generalization Principle

The universality of reason leads immediately to the *generalization principle*: a rational agent must believe that his/her reasons for acting are consistent with the assumption that all rational agents to whom the reasons apply could engage in the same actions (O’Neill, 2014; Wood, 1999).

As an example, suppose I see wristwatches on open display in a shop and steal one. My reasons for the theft are that I would like to have a new watch, and that I can get away with taking one.² At the same time, I cannot rationally believe that I would be able to get away with the theft if *everyone* stole watches when these reasons apply. The shop would install security measures to prevent theft, which is inconsistent with one of my reasons for stealing the watch. The theft therefore violates the generalization principle.

To give these ideas more precision, we express the action plan and generalization principle in the language of quantified modal logic. In so doing, we do not define a deductive system or propose formal semantics, as they are unnecessary for our project. We merely borrow logical notation in order to allow a more rigorous formulation and application of ethical principles.

The decision to steal a watch can be expressed in logical notation as follows. Define predicates

$$\begin{aligned}
 C_1(a) &= \text{Agent } a \text{ would like to possess an item on} \\
 &\quad \text{display in a shop.} \\
 C_2(a) &= \text{Agent } a \text{ can get away with stealing the item.} \\
 A_1(a) &= \text{Agent } a \text{ will steal the item.}
 \end{aligned}$$

²In practice, the reasons for theft are likely to be more complicated than this. I may be willing to steal partly because I believe the shop can easily withstand the loss, no employee will be disciplined or terminated due to the loss, I will not feel guilty afterward, and so forth. But for purposes of illustration we suppose there are only two reasons.

Because the agent's reasons are an essential part of moral assessment, we evaluate the agent's *action plan*, which states that the agent will take a certain action when certain reasons apply. In this case, the action plan is

$$(C_1(a) \wedge C_2(a)) \Rightarrow_a A_1(a) \quad (1)$$

Here \Rightarrow_a is not logical entailment but indicates that agent a regards $C_1(a)$ and $C_2(a)$ as justifying $A_1(a)$. The reasons in the action plan should be the most general set of conditions that the agent takes as justifying the action. Thus the action plan refers to an item in a shop rather than specifically to a watch, because the fact that it is a watch is not relevant to the justification; what matters is whether the agent wants the item and can get away with stealing it.

We can now state the generalization principle using quantified modal logic. Let $C(a) \Rightarrow_a A(a)$ be an action plan for agent a , where $C(a)$ is a conjunction of the reasons for taking action $A(a)$. The action plan is generalizable if and only if

$$\diamond_a P(\forall x(C(x) \Rightarrow_x A(x)) \wedge C(a) \wedge A(a)) \quad (2)$$

Here $P(S)$ means that it is physically possible for proposition S to be true, and $\diamond_a S$ means that a can rationally believe S . The proposition $\diamond_a S$ is equivalent to $\neg \square_a \neg S$, where $\square_a \neg S$ means that rationality requires require a to deny S .³ Thus (2) says that agent a can rationally believe that it is possible for everyone to have the same action plan as a , even while a 's reasons still apply and a takes the action.

Returning to the theft example, the condition (2) becomes the test proposition for action plan (2):

$$\diamond_a P(\forall x(C_1(x) \wedge C_2(x) \Rightarrow_x A_1(x)) \wedge C_1(a) \wedge C_2(a) \wedge A_1(a)) \quad (3)$$

This says that it is rational for a to believe that it is physically possible for the following to be true simultaneously: (a) everyone steals when the stated conditions apply, and (b) the conditions apply and a steals. Since (3) is false, action plan (1) is unethical.

The necessity of (3) for the rightness of action plan (1) is anchored in deontological theory, while the falsehood of (3) is a fact about the world. This fact might be inferred by collecting responses from shop owners about how they would react if theft were widespread. *Thus ethics and empirical VA*

³The operators \diamond and \square have a somewhat different interpretation here than in traditional epistemic and doxastic modal logics, but the identity $\diamond S \equiv \neg \square \neg S$ holds as usual.

work together in a very specific way: ethics tells us that the test proposition (3) must be true if the theft is to be ethical, and empirical VA provides evidence that bears on whether (3) is true.

An action plan in the autonomous vehicle domain might be

$$C_3(a) \Rightarrow_a A_2(a) \tag{4}$$

where

$C_3(a)$ = An ambulance under the control of agent a can reach its destination sooner by using siren and lights.

$A_2(a)$ = Agent a will direct an ambulance to use siren and lights.

Agent a is the ambulance driver, or in the case of an autonomous vehicle, the designer of the software that controls the ambulance. The generalization principle yields the test proposition

$$\diamond_a P \left(\forall x (C_3(x) \Rightarrow_y A_2(x)) \wedge C_3(a) \wedge A_2(a) \right) \tag{5}$$

This says that it is rational for agent a to believe that siren and lights could continue to hasten arrival if all ambulances used them for all trips, emergencies and otherwise. If empirical VA reveals that most drivers would ignore siren and lights if they were universally abused in this fashion, then we have evidence that (5) is false, in which case action plan (4) is unethical.

4.3 Maximizing Utility

Utilitarianism is normally understood as a *consequentialist* theory that evaluates an act by its actual consequences. Specifically, an act is ethical only if it maximizes total net expected utility across all who are affected. Yet the utilitarian principle can also be construed in a deontological fashion (Cummiskey, 1996), which allows it to be interpreted as requiring the agent to select actions that the agent can rationally believe will maximize utility. While utilitarians frequently view utility maximization as the sole ethical principle, it can be seen as an additional necessary condition for an ethical action. The other non-utilitarian principles remain in force because only actions that satisfy the other principles are considered options for maximizing utility.

In a deontological analysis, utility is not what people generally value but what the agent is rationally committed to valuing. The logic of means and ends requires that the agent regard some end as *intrinsically* valuable (such as happiness), and the universality of reason requires that it be seen

as valuable for any agent. A utilitarian believes this commits the agent to selecting actions that maximize the expected net sum of utility over everyone who is affected.⁴

The utilitarian principle can be formalized by requiring that a given action plan create at least as much utility as any other available action plan. Let $u(C(a), A(a))$ be a utility function that measures the total net expected utility of action $A(a)$ under conditions $C(a)$. Then an action plan $C(a) \Rightarrow_a A(a)$ satisfies the utilitarian principle only if agent a can rationally believe that action $A(a)$ creates at least as much utility as any ethical action that is available under the same circumstances. This can be written

$$\Diamond_a \forall A' \left(E(C(a), A'(a)) \rightarrow u(C(a), A(a)) \geq u(C(a), A'(a)) \right) \quad (6)$$

where A' ranges over actions. The predicate $E(C(a), A'(a))$ means that action $A'(a)$ is available for agent a under conditions $C(a)$, and that the action plan $C(a) \Rightarrow_a A'(a)$ is generalizable and respects autonomy.⁵ Note that we are now quantifying over predicates and have therefore moved into second-order logic.

Popular views about acceptable behavior frequently play a role in applications of the utilitarian principle. For example, in some parts of the world, drivers consider it wrong to enter a stream of moving traffic from a side street without waiting for a gap in the traffic. In other parts of the world this can be acceptable, because drivers in the main thoroughfare expect it and make allowances. Suppose driver a 's action plan is $(C_4(a) \wedge C_5(a)) \Rightarrow_a A_3(a)$, where

- $C_4(a)$ = Driver a wishes to enter a main thoroughfare.
- $C_5(a)$ = Driver a can enter a main thoroughfare by moving into the traffic without waiting for a gap.
- $A_3(a)$ = Driver a will move into traffic without waiting for a gap.

As before, driver a is the designer of the software if the vehicle is autonomous. Using (6), the driver's action plan maximizes utility only if the following test

⁴Alternatively, one might argue that maximizing the minimum utility over those affected (or achieving a lexicographic maximum) is the rational way to take everyone's utility into account, after the fashion of John Rawls's difference principle (Rawls, 1971). Or one might argue for some rational combination of utilitarian and equity objectives (Karsu and Morton, 2015; Hooker and Williams, 2012). However, for many practical applications, simple utility maximization appears to be a sufficiently close approximation to a "rational" choice, and to simplify exposition we assume so in this paper.

⁵For "respecting autonomy," see the next section.

proposition is true:

$$\diamond_a \forall A' \left(E(C_4(a), C_5(a), A'(a)) \rightarrow u(C_4(a), C_5(a), A_3(a)) \geq u(C_4(a), C_5(a), A'(a)) \right) \quad (7)$$

Suppose we wish to design driving policy in a context where pulling immediately into traffic is considered unacceptable. Then doing so is a dangerous move that no one is expecting, and an accident could result. Waiting for a gap in the traffic results in greater net expected utility, or formally, $u(C_4(a), C_5(a), A_3(a)) < u(C_4(a), C_5(a), A_4(a))$, where $A_4(a)$ is the action of moving into traffic after waiting for a gap. So (7) is false, and its falsehood can be inferred by collecting popular views about acceptable driving behavior. Observed preferences and values are therefore relevant to an ethical assessment, but they alone do not determine the assessment.

Again we have a clear demonstration of how ethical principles can combine with empirical VA. The utilitarian principle tells us that a particular action plan is ethical only if test proposition (7) is true, and empirical VA tells us whether (7) is true.

A similar approach can accommodate other situations in which popular expectations bear on ethical decisions. For example, it has been observed that people may expect different ethical norms to be followed by machine agents than by humans (Malle et al., 2015). This could affect generalizability as well as a utilitarian assessment, because there may be different implied promises or agreements concerning machines than humans. Yet again, expectations alone do not determine the ethical outcome.

4.4 Respect for Autonomy

A third ethical principle requires agents to respect the autonomy of other agents. Specifically, an agent should not adopt an action plan that the agent is rationally constrained to believe is inconsistent with an ethical action plan of another agent, without informed consent. Murder, enslavement, and inflicting serious injury are extreme examples of autonomy violations because they interfere with many ethical action plans. Coercion may or may not violate autonomy, depending on precisely how action plans are formulated.⁶

⁶A more adequate analysis leads to a principle of *joint* autonomy, according to which it is violation of autonomy to adopt an action plan that is mutually inconsistent with action plans of a set of other agents, when those other action plans are themselves mutually consistent. Joint autonomy addresses situations in which an action necessarily interferes with the action plan of some agent but no particular agent, as when someone throws a

The argument for respecting autonomy is basically as follows. Suppose I violate someone’s autonomy for certain reasons. That person could, at least conceivably, have the same reasons to violate my autonomy. This means that, due to the universality of reason, I am endorsing the violation of my own autonomy in such a case. This is a logical contradiction, because it implies that I am deciding not to do what I decide to do. To avoid contradicting myself, I must avoid interfering with other action plans.

To formulate an autonomy principle, we say that agent a ’s action plan $C(a) \Rightarrow_a A(a)$ is consistent with b ’s action plan $C'(b) \Rightarrow_b A'(b)$ when

$$\Diamond_a P(A(a) \wedge A'(b)) \vee \neg \Box_a P(C(a) \wedge C'(b)) \quad (8)$$

This says that agent a can rationally believe that the two actions are mutually consistent, or can rationally believe that the reasons for the actions are mutually inconsistent. The latter suffices to avoid inconsistency of the action plans, because if the reasons for them cannot both apply, the actions can never come into conflict.

As an example of how coercion need not violate autonomy, suppose agent b wishes to catch a bus and has decided to cross the street to a bus stop, provided no traffic is coming. The agent’s action plan is

$$(C_6(b) \wedge C_7(b) \wedge \neg C_8(b)) \Rightarrow_b A_5(b) \quad (9)$$

where

- $C_6(b)$ = Agent b wishes to catch a bus.
- $C_7(b)$ = There is a bus stop across the street from b .
- $C_8(b)$ = There are cars approaching b .
- $A_5(b)$ = Agent b will cross the street.

Agent a sees agent b begin to cross the street and forcibly pulls b out of the path of an oncoming car that b does not notice. Agent a ’s action plan is

$$(C_8(b) \wedge C_9(b)) \Rightarrow_a A_6(a, b) \quad (10)$$

where

- $C_9(b)$ = Agent b is about to cross the street.
- $A_6(a, b)$ = Agent a will prevent agent b from crossing the street.

bomb into a crowd. A general formulation of the joint autonomy principle in terms of modal operators is given in Hooker and Kim (2018). This and other complications are discussed in Hooker (2018).

Agent a does not violate agent b 's autonomy, even though there is coercion. Their action plans (9) and (10) are consistent with each other, because the condition (8) yields the test proposition

$$\diamond_a P(A_5(b) \wedge A_6(a, b)) \vee \neg \Box_a P(C_6(b) \wedge C_7(b) \wedge \neg C_8(b) \wedge C_8(b) \wedge C_9(b)) \quad (11)$$

This means that either (a) agent a can rationally believe that the two actions are consistent with each other, or (b) agent a can rationally believe that the antecedents of (9) and (10) are mutually inconsistent. As it happens, the two actions are obviously not consistent with each other, and so (a) is false. However, agent a can rationally believe that the antecedents of (9) and (10) are mutually inconsistent, because $C_8(b)$ and $\neg C_8(b)$ are contradictory. This means (b) is true, which implies that condition (11) is satisfied, and there is no violation of autonomy.

Again, this clearly distinguishes the roles of ethics and empirical observation in VA. Ethical reasoning tells us that the test proposition (11) must be true if autonomy is to be respected, whereas observation of the world tells us whether (11) is true.

In saying that coercion can be ethical, we do not imply that a violation of autonomy can be ethical. Coercion must be consistent with the coerced agent's action plan, as in the above example. Coercion can also be ethical when there is implied or informed consent, or when it is necessary to prevent unethical behavior, as in self-defense.⁷

To illustrate how autonomy may play a role in the ethics of driving, suppose that a pedestrian b dashes in front of a 's rapidly moving car. Driver a can slam on the brake and avoid impact with the pedestrian, but another driver c is following closely, and a sudden stop could cause a crash. The driver a must choose between two possible action plans:

$$(C_{10}(a, b) \wedge C_{11}(a, c)) \Rightarrow_a A_7(a) \quad (12)$$

$$(C_{10}(a, b) \wedge C_{11}(a, c)) \Rightarrow_a \neg A_7(a) \quad (13)$$

⁷Coercion can be ethical when there is informed consent to a risk of interference, because giving informed consent is equivalent to including the possibility of interference as one of the antecedents of the action plan. This occurs, for example, when a medical test subject gives consent with the knowledge that an experimental drug may cause illness, even though administering a drug that turns out to be harmful is a form of coercion. Interfering with an unethical action plan is no violation of autonomy because an unethical action plan is, strictly speaking, not an action plan due to the absence of a coherent set of reasons for undertaking it. An action plan is considered unethical in this context when it violates the generalization or utility principle, or interferes with an action plan that does not violate one of these principles, and so on recursively. Thus coercion is ethical in an act of self-defense, or to stop someone from unethically harming others.

where

$C_{10}(a, b)$ = Pedestrian b is dashing in front of a 's car.

$C_{11}(a, c)$ = Driver c is closely following a 's car.

$A_7(a)$ = Agent a will immediately slam on the brake.

Meanwhile, the pedestrian b has any number of action plans that are clearly inconsistent with death or serious injury. Let $C_{12}(b) \Rightarrow_b A_8(b)$ be one of them. Also driver c of the other car (there is only one occupant) has action plans that are inconsistent with an injury. We suppose that $C_{13}(c) \Rightarrow_c A_9(c)$ is one of them.

We first check whether hitting the brake, as in action plan (12), is inconsistent with the other driver's action plan $C_{13}(c) \Rightarrow_c A_9(c)$. The test proposition is

$$\diamond_a P(A_7(a) \wedge A_9(c)) \vee \neg \Box_a P(C_{10}(a, b) \wedge C_{11}(a, c) \wedge C_{13}(c)) \quad (14)$$

The first disjunct is clearly true, because a can rationally believe that it is *possible* that hitting the brake is consistent with avoiding a rear-end collision and therefore with any planned action $C_{13}(c) \Rightarrow_c A_9(c)$, even if this is improbable. So action plan (12) does not violate joint autonomy.

We now check whether a failure to hit the brake, as in action plan (13), is inconsistent with the pedestrian's action plan $C_{12}(b) \Rightarrow_b A_8(b)$. There is no violation of autonomy if

$$\diamond_a P(\neg A_7(a) \wedge A_8(b)) \vee \neg \Box_a P(C_{10}(a, b) \wedge C_{11}(a, c) \wedge C_{12}(b)) \quad (15)$$

The first disjunct of (15) is clearly false for b 's action plan $C_{12}(b) \Rightarrow_b A_8(b)$, because driver a cannot rationally believe that a failure to hit the brake is consistent with it. The second disjunct is likewise false, because driver a has no reason to believe that $C_{10}(a, b)$, $C_{11}(a, c)$ and $C_{12}(b)$ are mutually inconsistent. Thus (15) is false, and we have a violation of autonomy. The driver should therefore slam on the brake. There is no need to check the other ethical principles, because only one of the possible action plans satisfies the autonomy principle.

4.5 Implementation issues

While it is not our purpose to address engineering aspects of deontically-grounded VA, we can take note of some implementation issues that arise. The main implication of our proposal is that the portion of an AI system that makes ethically relevant decisions must be rule-based (i.e., an instance

of GOFAI) because it must consist of action plans. Fortunately, action plans have an if-then structure that is convenient for coding rules.

One can ask whether a rule-based system is adequate for the complexities of real-life decision making, but this is, of course, a problem that is not confined to deontically-based VA. We do not attempt here to judge the versatility of rule-based AI, but we note that it seems to be increasingly viewed as technically viable and even necessary due to the nontransparency of deep learning and support vector machines. Regarding autonomous vehicles, for example, Brandom (2018) states, “Many companies have shifted to rule-based AI, an older technique that lets engineers hard-code specific behaviors or logic into an otherwise self-directed system.” The technical community has ample experience at accurately coding and debugging huge rule-based systems. An ordinary (non-self-driving) automobile is already regulated by more than 100,000 lines of code. Ethics-based systems can evolve through several versions and be updated as necessary, as with any other type of complex software. Rule-based AI can also be combined with machine learning (Woźniak and Połap, 2020). Even in a pure ML system, it is possible to derive rules that approximate the directives generated by ML (Soares et al., 2020) and perhaps subject them to ethical evaluation.

The test propositions used to evaluate the ethical status of action plans need not appear in the AI rule base, and it is a further implementation decision whether to generate them automatically. This is fairly straightforward (less so for the utilitarian test), because the procedure for doing so can be clearly specified as shown above. Machine learning and other forms of empirical VA can then be used to evaluate the truth of the test propositions.

5 Conclusion

As AI inexorably enters everyday life, it takes a seat alongside human persons. AI’s increasing sophistication bestows power, and power begets responsibility. Humanity’s goal should be to invest machines with a moral sensitivity that mimics the human conscience. But conscience is dynamic rather than static, and adjusts ethical principles systematically to empirical observations. In this paper we have elaborated two challenges to AI moral reasoning that spring from the interrelation of facts and values. The first is a confusion that mistakenly identifies facts for values; the second is a confusion that misunderstands the process of moral reasoning. In addressing these challenges, we have identified instances of how and why AI can commit the naturalistic fallacy, of moving illicitly from “is’s” to “oughts,” and doing so oversimplifies

the process of moral reasoning. We have sketched, in response, a proposal for understanding moral reasoning in machines, one that highlights how deontological ethical principles can interact with factual states of affairs.

References

- Pieter Abbeel and Andrew Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the 21st International Conference on Machine Learning*. ACM, 2004.
- Raquel G. Alhama and Willem Zuidema. A review of computational models of basic rule learning: The neural-symbolic debate and beyond. *Psychonomic Bulletin & Review*, 26(4):1174–1194, 2019.
- Colin Allen. Calculated morality: Ethical computing in the limit. *Cognitive, emotive and ethical aspects of decision making and human action*, 1:19–23, 2002.
- Colin Allen, Gary Varner, and Jason Zinser. Prolegomena to any future artificial moral agent. *Journal of Experimental & Theoretical Artificial Intelligence*, 12(3):251–261, 2000.
- Colin Allen, Iva Smit, and Wendell Wallach. Artificial morality: Top-down, bottom-up, and hybrid approaches. *Ethics and Information Technology*, 7: 149–155, 2005.
- Colin Allen, Wendell Wallach, and Iva Smit. Why machine ethics? *IEEE Intelligent Systems*, 21(4):12–17, 2006.
- S. L. Anderson and M. Anderson. A prima facie duty approach to machine ethics. In M. Anderson and S. L. Anderson, editors, *Machine Ethics*, pages 476–492. Cambridge University Press, New York, 2011.
- Konstantine Arkoudas, Selmer Bringsjord, and Paul Bello. Toward ethical robots via mechanized deontic logic. In *AAAI Fall Symposium on Machine Ethics*, pages 17–23. The AAAI Press Menlo Park, CA, 2005.
- Thomas Arnold and Matthias Scheutz. The big red button is too late: An alternative model for the ethical evaluation of AI systems. *Ethics and Information Technology*, 20(1):59–69, 2018.
- Thomas Arnold, Daniel Kasenberg, and Matthias Scheutz. Value alignment or misalignment—What will keep systems accountable? In *3rd International Workshop on AI, Ethics, and Society*, 2017.
- Max H. Bazerman. Bounded ethicality in negotiations. *Negotiation and Conflict Management Research*, 4(1):8–11, 2011.

- Max H. Bazerman and Francesca Gino. Behavioral ethics: Toward a deeper understanding of moral judgment and dishonesty. *Annual Review of Law and Social Science*, 8:85–104, 2012.
- Max H. Bazerman and Don Moore. Is it time for auditor independence yet? *Accounting, Organizations and Society*, 36(4-5):310–312, 2011.
- Max H. Bazerman and Ovul Sezer. Bounded awareness: Implications for ethical decision making. *Organizational Behavior and Human Decision Processes*, 136:95–105, 2016.
- Max H. Bazerman and Ann E. Tenbrunsel. *Blind Spots: Why We Fail To Do What's Right and What To Do About It*. Princeton University Press, 2011.
- Selim Berker. The normative insignificance of neuroscience. *Philosophy & Public Affairs*, 37(4):293–329, 2009.
- Marianne Bertrand, Dolly Chugh, and Sendhil Mullainathan. Implicit discrimination. *American Economic Review*, 95(2):94–98, 2005.
- A. Bilgrami. *Self-Knowledge and Resentment*. Harvard University Press, Cambridge, MA, 1996.
- Russell Brandom. Self-driving cars are headed toward an AI roadblock. *The Verge*, 3 Jul 2018.
- Selmer Bringsjord. *What Robots Can and Can't Be*, volume 12. Springer Science & Business Media, 2013.
- Selmer Bringsjord. A 21st-century ethical hierarchy for robots and persons. In *A World with Robots*, pages 47–61. Springer, 2017.
- Selmer Bringsjord and Joshua Taylor. The divine-command approach to robot ethics. *Robot Ethics: The Ethical and Social Implications of Robotics*, pages 85–108, 2012.
- Selmer Bringsjord, Konstantine Arkoudas, and Paul Bello. Toward a general logicist methodology for engineering ethically correct robots. *IEEE Intelligent Systems*, 21(4):38–44, 2006.
- Cameron Buckner and James Garson. Connectionism. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2019 edition, 2019.

- Emanuelle Burton, Judy Goldsmith, Sven Koenig, Benjamin Kuipers, Nicholas Mattei, and Toby Walsh. Ethical considerations in artificial intelligence courses. *AI magazine*, 38(2):22–34, 2017.
- Paco Calvo and John Symons (eds.). *The Architecture of Cognition: Rethinking Fodor and Pylyshyn’s Systematicity Challenge*. MIT Press, 2014.
- Jean-Marie Chauvet. The 30-year cycle in the AI debate. *arXiv preprint arXiv:1810.04053*, 2018.
- Dolly Chugh, Max H. Bazerman, and Mahzarin R. Banaji. Bounded ethicality as a psychological barrier to recognizing conflicts of interest. In Don A. Moore, Daylian M. Cain, George Loewenstein, and Max H. Bazerman, editors, *Conflicts of Interest: Challenges and Solutions in Business, Law, Medicine, and Public Policy*, pages 74–95. Cambridge University Press, 2005.
- Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. Moral decision making frameworks for artificial intelligence. In *Proceedings of 31st AAAI Conference on Artificial Intelligence*, pages 4831–4835, 2017.
- David Gauthier. *Kantian Consequentialism*. Oxford University Press, 1996.
- David Danks. Learning. In Frankish Keith and William M. Ramsey, editors, *Cambridge Handbook to Artificial Intelligence*, pages 151–167. Cambridge University Press, 2014.
- Thomas Donaldson. When integration fails: The logic of prescription and description in business ethics. *Business Ethics Quarterly*, 4(2):157–169, 1994.
- Thomas Donaldson. The epistemic fault line in corporate governance. *Academy of Management Review*, 37(2):256–271, 2012.
- Miguel Egler and Lewis D. Ross. Philosophical expertise under the microscope. *Synthese*, 197(3):1077–1098, 2020.
- Jerry A. Fodor and Zenon W. Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2):3–71, 1988.

- Jean-Gabriel Ganascia. Modelling ethical rules of lying with answer set programming. *Ethics and information technology*, 9(1):39–47, 2007.
- Hector Geffner. Model-free, model-based, and general intelligence. *arXiv preprint arXiv:1806.02308*, 2018.
- Naveen Sundar Govindarajulu and Selmer Bringsjord. On automating the doctrine of double effect. *arXiv preprint arXiv:1703.08922*, 2017.
- Alexander R. Green, Dana R. Carney, Daniel J. Pallin, Long H. Ngo, Kristal L. Raymond, Lisa I. Iezzoni, and Mahzarin R. Banaji. Implicit bias among physicians and its prediction of thrombolysis decisions for black and white patients. *Journal of General Internal Medicine*, 22(9):1231–1238, 2007.
- Anthony G. Greenwald, T. Andrew Poehlman, Eric Luis Uhlmann, and Mahzarin R. Banaji. Understanding and using the implicit association test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, 97(1):17, 2009.
- Richard Mervyn Hare. *The Language of Morals*. Oxford Paperbacks, 1991.
- John Haugeland. *Artificial Intelligence: The Very Idea*. MIT Press, 1985.
- J. N. Hooker and Tae Wan Kim. Truly autonomous machines are ethical. *AI Magazine*, 40(4), winter 2019.
- J. N. Hooker and Tae Wan N. Kim. Toward non-intuition-based machine and artificial intelligence ethics: A deontological approach based on modal logic. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '18, pages 130–136, 2018.
- J. N. Hooker and H. P. Williams. Combining equity and utilitarianism in a mathematical programming model. *Management Science*, 58(9):1682–1693, 2012.
- John Hooker. *Taking Ethics Seriously: Why Ethics Is an Essential Tool for the Modern Workplace*. Taylor & Francis, 2018.
- Daniel Kahneman. *Thinking, Fast and Slow*. Macmillan, 2011.
- Immanuel Kant. *Grundlegung zur Metaphysik der Sitten (Foundations of the Metaphysics of Morals)*, volume 4 of *Königlichen Preußischen Akademie der Wissenschaften: Kants gesammelte Schriften*. Georg Reimer (1900), Berlin, 1785.

- Ö. Karsu and A. Morton. Inequity averse optimization in operational research. *European Journal of Operational Research*, 245:343–359, 2015.
- C. M. Korsgaard. *The Sources of Normativity*. Cambridge University Press, Cambridge, 1996.
- Brenden Lake and Marco Baroni. Still not systematic after all these years: On the compositional skills of sequence-to-sequence recurrent networks. *ArXiv preprint 1711.00350*, 2018.
- Patrick Lin, Keith Abney, and George A. Bekey (eds.). *Robot Ethics: The Ethical and Social Implications of Robotics*. MIT press, 2011.
- Felix Lindner and Martin Mose Bentzen. A formalization of Kant’s second formulation of the categorical imperative. *arXiv preprint 1801.03160*, 2018.
- Felix Lindner, Robert Mattmüller, and Bernhard Nebel. Evaluation of the moral permissibility of action plans. *Artificial Intelligence*, 2020.
- Bertram F. Malle, Matthias Scheutz, Thomas Arnold, John Voiklis, and Corey Cusimano. Sacrifice one for the good of many?: People apply different moral norms to human and robot agents. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 117–124, 2015.
- Gary Marcus. *The Algebraic Mind*. MIT Press, 2001.
- Gary Marcus. Deep learning: A critical appraisal. *arXiv preprint 1801.00631*, 2018.
- John McDowell. Two sorts of naturalism. In Hursthouse R., Lawrence G., and Quinn W., editors, *Virtues and Reasons: Philippa Foot and Moral Theory: Essays in Honour of Philippa Foot*. Clarendon Press, 1995.
- Don A. Moore, Philip E. Tetlock, Lloyd Tanlu, and Max H. Bazerman. Conflicts of interest and the case of auditor independence: Moral seduction and strategic issue cycling. *Academy of Management Review*, 31(1):10–29, 2006.
- Don A. Moore, Lloyd Tanlu, and Max H. Bazerman. Conflict of interest and the intrusion of bias. *Judgment and Decision Making*, 5(1):37, 2010.
- Thomas Nagel. Ethics without biology. In Thomas Nagel, editor, *Mortal Questions*, pages 142–46. Cambridge University Press, 1979.

- Thomas Nagel. *The View from Nowhere*. Oxford University Press, Oxford, 1986.
- D. K. Nelkin. Two standpoints and the belief in freedom. *Journal of Philosophy*, 97:564–576, 2000.
- Andrew Y. Ng and Stuart J. Russell. Algorithms for inverse reinforcement learning. In *Proceedings 17th International Conference on Machine Learning*, pages 663–670. Morgan Kaufmann, 2000.
- Onora O’Neill. *Acting on Principle: An Essay on Kantian Ethics*, 2nd ed. Cambridge University Press, Cambridge, UK, 2014.
- Derek Parfit. *On What Matters*, volume 1. Oxford University Press, 2011.
- Charles Pigden. Hume on *is* and *ought*: Logic, promises, and the Duke of Wellington. In *The Oxford Handbook of Hume*, pages 401–415. Oxford University Press, 2016.
- Alan Prince and Steven Pinker. Rules and connections in human language. *Trends in Neurosciences*, 11(5):195–202, 1988.
- James Rachels. *Created from Animals*. Oxford University Press, 1990.
- John Rawls. *A Theory of Justice*. Harvard University Press, 1971.
- Michael Rescorla. The computational theory of mind. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2020 edition, 2020.
- Francesca Rossi and Andrea Loreggia. Preferences and ethical priorities: Thinking fast and slow in AI. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 3–4, 2019.
- Laurie A. Rudman and Richard D. Ashmore. Discrimination and the implicit association test. *Group Processes & Intergroup Relations*, 10(3):359–372, 2007.
- Stuart Russell, Daniel Dewey, and Max Tegmark. Research priorities for robust and beneficial artificial intelligence. *AI Magazine*, 36(4):105–114, 2015.
- Matthias Scheutz. Computationalism—The next generation. In *Computationalism: New Directions*, pages 1–21. 2002.

- Matthias Scheutz and Thomas Arnold. Are we ready for sex robots? In *11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 351–358. IEEE, 2016.
- Eric Schwitzgebel and Fiery Cushman. Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language*, 27(2):135–153, 2012.
- Eric Schwitzgebel and Fiery Cushman. Philosophers’ biased judgments persist despite training, expertise and reflection. *Cognition*, 141:127–137, 2015.
- Ovul Sezer, Francesca Gino, and Max H. Bazerman. Ethical blind spots: Explaining unintentional unethical behavior. *Current Opinion in Psychology*, 6:77–81, 2015.
- Daniel J. Singer. Mind the is-ought gap. *The Journal of Philosophy*, 112(4): 193–210, 2015.
- Jeffery Smith and Wim Dubbink. Understanding the role of moral principles in business ethics: A Kantian perspective. *Business Ethics Quarterly*, 21(2):205–231, 2011.
- Paul Smolensky. On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11(1):1–23, 1988.
- E. A. Soares, P. P. Angelov, and B. Costa. Explaining deep learning models through rule-based approximation and visualization. *IEEE Transactions on Fuzzy Systems*, published online 3 Jun 2020.
- Ann E Tenbrunsel. Commentary: Bounded ethicality and conflicts of interest. In Don A. Moore, Daylian M. Cain, George Loewenstein, and Max H. Bazerman, editors, *Conflicts of Interest: Challenges and Solutions in Business, Law, Medicine, and Public Policy*, page 96. Cambridge University Press, 2005.
- Kevin Tobia, Wesley Buckwalter, and Stephen Stich. Moral intuitions: Are philosophers experts? *Philosophical Psychology*, 26(5):629–638, 2013.
- Wendell Wallach and Colin Allen. *Moral Machines: Teaching Robots Right from Wrong*. Oxford University Press, 2008.
- Wendell Wallach, Colin Allen, and Iva Smit. Machine morality: Bottom-up and top-down approaches for modelling human moral faculties. *AI & Society*, 22(4):565–582, 2008.

- Toby Walsh, Neil Levy, Genevieve Bell, Anthony Elliott, James Maclaurin, Iven Mareels, and Fiona Wood. *The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve our Wellbeing*. Australian Council of Learned Academies, 2019.
- Alex Wiegmann, Joachim Horvath, and Karina Meyer. Intuitive expertise and irrelevant options. *Oxford Studies in Experimental Philosophy Volume 3*, 3:275, 2020.
- M. J. Wolf, K. Miller, and F. S. Grodzinsky. Why we should have seen that coming: Comments on Microsoft’s Tay “experiment,” and wider implications. *SIGCAS Computers and Society*, 47(3):54–64, 2017.
- Allen W. Wood. *Kant’s Ethical Thought*. Cambridge University Press, 1999.
- Jack Woods and Barry Maguire. Model theory, Hume’s dictum, and the priority of ethical theory. *Ergo*, 14(4), 2017.
- Marcin Woźniak and Dawid Połap. Intelligent home systems for ubiquitous user support by using neural networks and rule-based approach. *IEEE Transactions on Industrial Informatics*, 16(4):2651–2658, 2020.
- Ting Zhang, Pinar O. Fletcher, Francesca Gino, and Max H. Bazerman. Reducing bounded ethicality: How to help individuals notice and avoid unethical behavior. *Organizational Dynamics*, pages 310–317, 2015.