

MORAL DILEMMAS FOR MORAL MACHINES

TRAVIS LACROIX

ABSTRACT. Autonomous systems are being developed and deployed in situations that may require some degree of ethical decision-making ability. As a result, research in machine ethics has proliferated in recent years. This work has included using moral dilemmas as validation mechanisms for implementing decision-making algorithms in ethically-loaded situations. Using trolley-style problems in the context of autonomous vehicles as a case study, I argue (1) that this is a misapplication of philosophical thought experiments because (2) it fails to appreciate the purpose of moral dilemmas, and (3) this has potentially catastrophic consequences; however, (4) there are uses of moral dilemmas in machine ethics that are appropriate and the novel situations that arise in a machine-learning context can shed some light on philosophical work in ethics.

Keywords — AI Ethics, Moral Dilemmas, Artificial Moral Agency, Thought Experiments

1. INTRODUCTION

Increasingly, autonomous systems are being developed and deployed in situations that may require some degree of ethical decision-making ability. Some well-discussed examples include autonomous weapons for warfare (Arkin, 2008a,b; Krishnan, 2009; Tonkens, 2012; Hellström, 2013; Asaro, 2020); professional service robots for health-care and elderly care (Anderson et al., 2006; Anderson and Anderson, 2008; Sharkey and Sharkey, 2012; Conti et al., 2017); sex robots for therapy or personal pleasure (Eichenberg et al., 2019; Headleand et al., 2020; Döring et al., 2020); and self-driving vehicles for transportation (Bhargava and Kim, 2017; Sommaggio and Marchiori, 2018; Evans et al., 2020).

In the early days of machine learning (ML), researchers could focus on the fundamental aspects of their work without much concern for social or ethical consequences since these systems were relatively encapsulated within the confines of

DEPARTMENT OF PHILOSOPHY DALHOUSIE UNIVERSITY

E-mail address: tlacroix@dal.ca.

Date: Draft of March 11, 2022. Please cite official version, <https://doi.org/10.1007/s43681-022-00134-y>.

the research lab. However, [Luccioni and Bengio \(2019\)](#) highlight that these algorithms are increasingly being deployed in society. This is due, in part, to the promise of the unprecedented economic impacts of ML applications ([Bughin et al., 2019](#); [Szczepański, 2019](#); [Russell, 2019](#)). As a result, research in machine ethics—including fundamental questions surrounding the very nature and possibility of artificial moral agency—has proliferated in recent years.¹ This work has included using moral dilemmas (i.e., philosophical thought experiments) as validation mechanisms for implementing decision-making algorithms in ethically-loaded situations.

This paper aims to describe the use of philosophical thought experiments in the context of machine ethics research and explain how these experiments are misused in this field. As I argue, this misapplication comes from an apparent misunderstanding of what morally charged thought experiments from philosophy are supposed to accomplish. I conclude by describing what philosophical thought experiments are useful for in the context of ML and addressing some meta-ethical worries. I also describe how the novel situations that arise from the possibility of autonomous agents can shed some further light on philosophical work in ethics.

As a concrete example, I will focus on trolley-style problems applied to hypothetical scenarios that may be faced by autonomous vehicles since this is perhaps the most prevalent (mis-)use of a philosophical thought experiment in the context of artificial intelligence systems; however, I will also gesture toward other examples when applicable, to not give the (false) impression that this is a relatively isolated case.

2. MORAL MACHINES

In this section, I discuss how philosophical thought experiments—particularly, moral dilemmas—are used in machine learning to benchmark the ‘ethical’ performance of new algorithms. As a case study, I begin by providing some technical

¹So too have attempts to codify principles for ethical AI research, though largely to little effect. See [Jobin et al. \(2019\)](#) for a recent survey; see also [LaCroix and Mohseni \(2020\)](#) for a discussion of the efficacy of such proposals.

background on how autonomous vehicles work, which makes salient several possible problems arising in situations where a decision may need to be rendered that has potential moral weight (2.1). This situation gives rise to the appropriation of trolley-style problems in the context of autonomous vehicles needing to ‘choose’ between two undesirable alternatives (2.2). I then describe a well-known use case of trolley-style problems in machine ethics: the *Moral Machine Experiment* (2.3). This section concludes with a discussion of moral dilemmas more generally, highlighting how they are taken in the machine learning literature as *benchmarks* for determining whether an algorithm ‘acts ethically’ (2.4). In the subsequent section, I argue that this view is mistaken.

2.1. Autonomous Vehicles. The Society for Automotive Engineers defines six levels of automation, ranging from (0) no automation, where the driver performs all driving tasks, to high (4) or full (5) automation, where the vehicle is capable of autonomously performing all driving functions under certain/all conditions. Most vehicles on the road today are classed as level 0 or 1: they are controlled by humans but may have some driver-assistance capabilities, such as adaptive cruise control. However, several vehicles on the market from several different manufacturers fall under level 2 or 3—partial and conditional automation, respectively. For example, Tesla’s level-2 autopilot function partially automates the vehicle, but a human driver’s attention is still legally required at all times. Honda was the first company to have a vehicle classed with the level-3 designation, although this model has yet to be mass-produced—as of September 2021, the public sale of this vehicle was limited to 100 units in Japan. Predictions vary widely as to when fully-autonomous vehicles will be available for private use, which is consistent with the long history of overestimating the near-future abilities of AI systems.²

²For example, the goal of the Dartmouth Summer Research Project on Artificial Intelligence, held in 1956 and organised by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon, was stated as follows:

The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts,

Machine learning algorithms for autonomous vehicles must continuously render the surrounding environment, in addition to predicting possible changes to that environment moving forward through time and reacting appropriately to those changes. This ability involves a suite of applications, including object detection, recognition, localisation, and prediction (of movement). For example, many autonomous vehicles utilise RADAR (radio detection and ranging) and LIDAR (light detection and ranging) sensors, in addition to video cameras (to measure distance, detect road edges, and identify lane markings) and ultrasonic sensors (to detect curbs and other vehicles). These data may be fed into (typically several) deep neural networks and processed in real-time.³

The advent of artificial intelligence systems highlights how difficult it is to perform tasks that humans take for granted. For example, the sensors of an AV must detect when a pedestrian is waiting at a crosswalk, recognise that it is a pedestrian, predict whether the pedestrian will step out into the road, and respond appropriately—i.e., by slowing down or stopping. And, strange things can happen when the system is presented with examples that it has not yet encountered: in 2018, a self-driving vehicle in Tempe, AZ apparently alternated between classifying a pedestrian, Elaine Herzberg, who was walking her bicycle in the street, as ‘vehicle’, ‘person’, and ‘other object’. The result was that the vehicle struck and killed Herzberg (Wakabayashi, 2018).

solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

Of course, these are all still open problems today (Russell, 2019).

³The individual tasks—detection, recognition, localisation, prediction, action—can be accomplished using several different methods—including regression, pattern recognition, clustering, and decision matrices, often involving a plethora of state-of-the-art machine-learning techniques. For example, support vector machines and principal component analysis may be used for pattern recognition; K-means clustering may be used to identify data in low-resolution images; and gradient-boosting may be used for decision-making, depending on confidence levels for detection, classification, and prediction. Advances are continually being made in this field, with some focus on end-to-end learning; for example, Bojarski et al. (2016); Yang et al. (2017) employ convolutional neural networks to train their vehicles without requiring the highly complex suite of algorithms used in traditional methods. Kuefler et al. (2017) use Generative Adversarial Networks to train their system by mimicking human behaviour.

Part of the difficulty arises from the training data. In Herzberg’s case, the system could not recognise a human walking with a bike as two separate things needing classification (among many other things that went wrong). In addition, the software used in this particular instance did not include considerations for jaywalking pedestrians. If training examples always include crosswalks, the system may pick up on these underlying regularities instead of the intended target. For example, suppose an image-classification algorithm has only ever seen red apples. In that case, it might misclassify a green apple as a pear because its ‘concept’ of APPLE depends (too) heavily on some spurious regularities in the examples it has seen (Christian, 2020).

Let’s suppose that these problems are surmountable and fully autonomous vehicles are achievable in the foreseeable future.⁴ As autonomous vehicles become more prevalent on society’s roads, it is supposed that it will become increasingly likely that an individual vehicle will need to be programmed to make decisions in situations that carry significant moral weight. Practically, this is a difficult problem. Crashes are relatively rare in terms of the data that a machine-learning system might receive; therefore, the system may not ‘know’ how to respond, because of a lack of training data. Low-probability, but high-*risk*, events pose particular challenges for machine learning methods that depend upon the system seeing many examples in order to learn. This is true even when there is an objectively correct answer to the problem; however, in morally-charged situations, there may not be obviously correct answers on which to train the model, as I will discuss below.

2.2. Trolley-Style Problems. In some (perhaps exceedingly rare) circumstances, an autonomous vehicle may face a situation that can be classified as a trolley-style problem. As is well-known in philosophy, the trolley problem is a set of ethical dilemmas wherein a subject must choose between some set of options involving (typically) human lives. The problem was first introduced in Foot (1967) in a discussion of abortion and the doctrine of double-effect. This initial problem was expanded

⁴In fact, full automation is not necessary for the problems described to arise since human reaction time will not be useful in split-second moral decisions.

upon and later analysed in much more detail by Thomson (1976) and Unger (1996). I refer to these as trolley-*style* problems because, although they have the same basic structure as a trolley problem, their scope has been expanded well beyond the original philosophical context.

Despite the rarity of trolley-style problems in the real world, the extent to which they are possible implies that the machine will need to ‘know’ how to react. As a result, the advent of autonomous vehicles has re-invigorated interest within and without philosophy on the subject of trolley problems; however, as noted in the introduction, this is but one salient example of a more general interest in moral dilemmas.⁵

We suppose that an autonomous vehicle is about to crash and has no trajectory to save everyone. Is it better, for example, to hit a group of pedestrians on the road or swerve into a barrier, killing the driver? When harm is possible or inevitable, the vehicle will need to make a decision, which means that it needs to have been programmed or trained to be capable of making a decision. And, this is true regardless of how rare the circumstances might be in practice. In response to these facts, Awad et al. (2018) have noted that it will be necessary to gauge social expectations about how to divide the risk of harm between the different stakeholders on the road. Their response to this is the *Moral Machine Experiment*.

2.3. The Moral Machine Experiment. The Moral Machine Experiment (Awad et al., 2018) is a multilingual online ‘game’ for gathering human perspectives on moral dilemmas—specifically, trolley-style problems in the context of autonomous

⁵In the last few years alone, there have been dozens of articles that refer to Philippa Foot’s 1967 paper in the context of autonomous vehicles; see, for example, Allen et al. (2011); Wallach and Allen (2009); Pereira and Saptawijaya (2015, 2011); Berreby et al. (2015); Danielson (2015); Lin (2015); Malle et al. (2015); Saptawijaya and Pereira (2015, 2016); Bentzen (2016); Bhargava and Kim (2017); Casey (2017); Cointe et al. (2017); Greene (2017); Lindner et al. (2017); Santoni de Sio (2017); Welsh (2017); Wintersberger et al. (2017); Bjørgen et al. (2018); Grinbaum (2018); Misselhorn (2018); Pardo (2018); Sommaggio and Marchiori (2018); Baum et al. (2019); Cunneen et al. (2019); Krylov et al. (2019); Sans and Casacuberta (2019); Wright (2019); Agrawal et al. (2020); Awad et al. (2020); Banks (2021); Bauer (2020); Etienne (2020); Gordon (2020); Harris (2020); Lindner et al. (2020); Nallur (2020). And, several more articles that discuss trolley problems without citing Foot; e.g., Bonnefon et al. (2016); Etzioni and Etzioni (2017); Lim and Taeihagh (2019); Evans et al. (2020); or, which appear to reinvent the trolley problem (without citing Foot); e.g., Keeling (2018).

vehicles. By the time of publication, the Moral Machine Experiment had collected nearly 40 million data points from around the world. Individuals who participated were shown many unavoidable accident scenarios with binary outcomes—i.e., trolley-style problems—and were prompted to choose the scenario they prefer. These include ‘sparing humans (versus pets), staying on course (versus swerving), sparing passengers (versus pedestrians), sparing more lives (versus fewer lives), sparing men (versus women), sparing the young (versus the elderly), sparing pedestrians who cross legally (versus jaywalking), sparing the fit (versus the less fit), and sparing those with higher social status (versus lower social status)’ (60). Some scenarios include other ‘characters’, such as criminals, pregnant women, or doctors.⁶ Globally, they find that individuals tend to prioritise humans over animals, many humans over fewer, and younger humans over older.

According to Edmond Awad—one of the paper’s coauthors—the original purpose of the Moral Machine Experiment was supposed to be purely descriptive, highlighting *people’s* preferences in ethical decisions (Vincent, 2018). However, the first, second, and last authors (Awad, Dsouza, Rahwan) of the original paper, citing the results of the *Moral Machine Experiment*, published an article a month later which proposes a ‘voting-based system for ethical decision making’ (Noothigattu et al., 2018). They suggest that the Moral Machine Experiment data can be used to automate decisions, ‘even in the absence of such ground-truth principles, by aggregating people’s opinions on ethical dilemmas’ (4). This statement takes the descriptive project into the realm of normative ethics by suggesting that the Moral Machine Experiment data can serve as a validation mechanism for whether an algorithm acts ‘morally’.

2.4. Benchmarking Ethical Decisions. Several authors have proposed algorithms for moral decision making in autonomous vehicles, and there are intrinsic reasons why we might want AI systems to be capable of acting ethically. However, for-profit corporations have additional incentives for designing ‘ethical’ AI since

⁶There are obvious and perhaps pressing philosophical questions that arise concerning some of these classifications, but we will put those aside for now.

humans (i.e., *consumers*) will likely be more trusting of an autonomous agent (i.e., *products*) if it is known to possess a set of moral principles intended to constrain and guide its behaviour (Bonnefon et al., 2016). The question then arises how we are supposed to know whether the decision chosen by the system is ‘in fact’ moral or not—i.e., *how* ethical are the decisions made by the algorithm?

Benchmarking is a way of evaluating and comparing new methods in ML for performance on a particular dataset (Olson et al., 2017). Following Raji et al. (2021), we can understand a *benchmark*, for the purposes of this paper, as a dataset plus a metric for measuring the performance of a particular model on a specific task. For example, suppose that the current state of the art of image classification on ImageNet for top-1 accuracy is 85%. In this case, 85% top-1 accuracy is a benchmark—namely, an objective measure of how well one’s algorithm performs on a particular dataset for image classification. So, if a method performs worse than this benchmark, we know that something has gone awry. However, if the new method performs *better* than this benchmark, it is the best-performing algorithm to date (again, modulo efficiency, the volume of training data, etc.). By and large, state-of-the-art progress on certain benchmarks is typically taken to indicate progress on a particular task or set of tasks (Raji et al., 2021).

Moral dilemmas have been appealed to as benchmarks for checking whether an algorithm makes the ‘right’ decision in machine ethics. The most popular method for *evaluating* whether an artificial system behaves ethically is by evaluating its performance on ethical dilemmas (Nallur, 2020). In the case of autonomous vehicles, the most common dilemma that is appealed to is the trolley problem—due, in no small part, to its use in the Moral Machine Experiment. Using the Moral Machine Experiment as a benchmark can be understood in the following way. The dataset is the survey data that was collected by the experimenters—i.e., which of the binary outcomes is preferred by participants, on average. If human agents strongly prefer sparing more lives to fewer, then researchers might conclude that the ‘right’ decision for their algorithm to make is the one that reflects this sociological fact. Thus, the

metric would measure how close the algorithm’s decision is to the aggregate survey data.⁷ Note, then, that *survey data* is being used as a proxy for *moral facts*.

As previously mentioned, the notion of using moral dilemmas as a benchmark for machine ethics extends well beyond the particular case of trolley-style problems and autonomous vehicles. [Lourie et al. \(2020\)](#) introduce a dataset of ethical dilemmas which they say ‘enables models to learn basic ethical understanding’. However, as with the Moral Machine Experiment, it is important to note that their metric for measuring performance can only measure how humans annotate the dilemma (i.e., whether a response is ethical or not) *on average*.

In the more general case of moral dilemmas as benchmarks, [Bonnemains et al. \(2018\)](#) explicitly reason as follows: since classic moral dilemmas have already been used as a basis for ethical reasoning, ‘it seems legitimate to use some of them as a starting point for designing an automated ethical judgement on decisions’ (43). [Björger et al. \(2018\)](#) argue that *certain* types of ethical dilemmas ‘can be used as benchmarks for estimating the ethical performance of an autonomous system’ (23). [Kim et al. \(2018\)](#) construct a hierarchical Bayesian model for inferring individual and shared moral values from sparse and noisy data. Then they *evaluate* this approach by comparing their results with data from the Moral Machine Experiment. [Cunneen et al. \(2019\)](#) suggest that the use of trolley-style problems as an elucidatory tool is a necessary *precedent* (i.e., is necessarily prior) to focusing AI applications on moral theories. And, [Sütfeld et al. \(2017\)](#) suggest that models of ethics (specifically for autonomous vehicles) should *aim* to match human decisions made in the same context.

So, drawing normative consequences (an ‘ought’) from some descriptive matters of fact (an ‘is’) is not an unusual move in the field of machine ethics—particularly

⁷Of course, this approach raises all the usual problems of biased data, insofar as certain individuals are going to be overrepresented—i.e., those individuals from countries who have easy access to the internet. [Falbo and LaCroix \(2021\)](#) argue that these considerations may exacerbate structural inequalities and mechanisms of oppression—although, they also note that ‘more data’ is not necessarily going to fix that, since the data are reflective of extant inequalities in society. However, it is crucial to note that the thing being measured in this case is *not* how ethical the decision is, but how closely the decision accords with the opinions of humans, on average.

in the case of autonomous vehicles and trolley-style moral dilemmas. Besides being logically unsound, Etienne (2020) argues that this type of project is a dangerous basis for machine ethics, insofar as Awad et al. (2018); Noothigattu et al. (2018) lean on concepts of *social acceptability*, rather than, e.g., *fairness* or *rightness*; furthermore, individual opinions about these dilemmas may be highly volatile over time (Henrich et al., 2001; House et al., 2013; Blake et al., 2015).

However, setting aside some of the logical and practical problems that using moral dilemmas as benchmarks entails, there is a sense in which these authors fail to appreciate the role of moral dilemmas in philosophical thought. In the next section, I provide a description of the *purpose* of philosophical thought experiments, including moral dilemmas, to show why using moral dilemmas for verification tools involves a category mistake. In the conclusion, I discuss how this sets a dangerous precedent in the design of ‘ethical’ AI systems.

3. PHILOSOPHICAL THOUGHT EXPERIMENTS

As mentioned, moral dilemmas (a specific type of philosophical thought experiment), have been used as verification tools for moral decision-making in AI systems. In the case of the trolley-style problems posed by the Moral Machine Experiment (a specific type of moral dilemma), datasets—consisting of human responses to binary hypotheticals—and a measure of the system’s predictive accuracy with respect to those responses (on average) are taken to provide a benchmark for determining how ‘ethically’ the system acts. Here, I suggest that this is a mistake because, among other things, it misses the *point* of philosophical thought experiments.

Thought experiments, historically, have served an important role in philosophical discourse—especially in the philosophy of science and metaphilosophy.⁸ According to Kuhn (1977), thought experiments (generally) could make salient the failure of the world to conform to prior expectations about the way the world is. Similarly,

⁸There is a *vast* metaphilosophical literature on the role and purpose of thought experiments in philosophy. I do not have the space to delve into any adequate detail here; however, see Brown and Fehige (2019) for an overview. See also Asikainen and Hirvonen (2014) for a discussion of thought experiments in the context of science.

thought experiments might elucidate particular ways in which a theory—i.e., one that is based on said prior expectations—might be revised to better conform with the facts about the world.

Following Brun (2018), we can give the following simple schema for a thought experiment:

- (1) A scenario and a question are introduced;
- (2) The experimenter elucidates the scenario and arrives at some result;
- (3) A conclusion is drawn about some target(s).

A quick gloss on the *purpose* of thought experiments is that they serve *primarily* as intuition pumps (Dennett, 1980, 1991, 2013). They may be used, for example, to elicit normative intuitions, to justify counterfactual claims (also relying on intuitions), to explore logical relationships among philosophical theses, among others (Mayo-Wilson and Zollman, 2020). As Dennett (1980) describes intuition pumps, they are typically used for *provoking* ‘a family of intuitions by producing variations on a basic thought experiment’ (429). Importantly, philosophical thought experiments (as intuition pumps), should not be understood as ‘an engine of discovery, but a persuader or pedagogical tool—a way of getting people to see things your way’ (429).

So, moral dilemmas, like the trolley problem, provide some basic structure. A comparison of cases elucidating apparently incompatible or inconsistent reactions is supposed to shed light on some (morally) *salient* differences between the cases. This, in turn allows us to theorise about possible or plausible explanations for those differences. However, this too depends on individual intuitions to some extent: responses to moral dilemmas vary widely across societies and time periods (Henrich et al., 2001; House et al., 2013; Blake et al., 2015). Trolley problems, specifically, have figured heavily in empirical research in neuroscience and psychology, and, again, human responses to these scenarios are highly dependent on external features.⁹

⁹See, for example, Greene et al. (2001, 2004, 2008); Nichols and Mallon (2005); Cushman et al. (2006); Schaich Borg et al. (2006); Ciaramelli et al. (2007); Hauser et al. (2007); Koenigs et al. (2007); Waldmann and Dieterich (2007); Moore et al. (2008).

A moral dilemma like the trolley problem pumps intuitions about why it may be permissible to perform an intentional action despite its foreseen and undesirable consequences. Intuitively, it seems like it may be permissible to pull a switch to divert the trolley from one track to another, as in the original trolley problem, despite that this would lead to the (foreseeable) death of one individual, because five people would be saved as a result.¹⁰ At the same time, it seems impermissible to actively kill someone—say, by pushing them on the track to block the trolley—despite that this would have an identical outcome: five lives saved at the expense of one.

Thus, the dilemma gives rise to the question: Why might it *seem* morally permissible to act in one case but not in the other. Foot’s proposed *explanation* for divergent intuitions in these two cases is that there is differential import between the *positive* and *negative* duties one has—in particular, negative rights (and the ‘duties’ which follow from them) typically outweigh positive rights. In this case, the target of analysis is actually about the ethics of abortion—the ethical issue, it should go without saying, is not about trolleys. The thought experiment is useful because people are less likely to carry pre-theoretic baggage about trolleys than about abortions. Therefore, the trolley problem gets at the core of the issue in applied ethics while abstracting away from the moral loadedness of the actual target.

This applies to philosophical thought experiments more generally, not just moral dilemmas. Consider a famous thought experiment: Gettier cases from epistemology (Gettier, 1963). In a Gettier case, an individual clearly has a true and justified belief about a proposition, p , but it is not obvious that they *know* p . This is supposed to show that justified true belief cannot be sufficient for knowledge. But, for this thought experiment to work, it must be that in a Gettier case, the individual in question has a justified true belief that p and *in fact* fails to know that p . However, knowledge is the very thing that epistemologists are trying to define—in this case, by deference to some set of (individually) necessary and (jointly) sufficient conditions.

¹⁰Surveys have shown that a vast majority would choose to pull the switch in this scenario Navarrete et al. (2012); Bourget and Chalmers (2014).

Thus, if knowledge is a concept that requires analysis, the Gettier case cannot show that justified true belief is not knowledge since we do not know, *a priori*, what knowledge is. These examples depend upon the reader sharing Gettier's intuition that there is no knowledge in either of these cases. As such, the success of the Gettier cases is more of a sociological success than an epistemic one—what it shows is that *many philosophers share the intuition that there is no knowledge in Gettier cases.*

To make this point clearer, consider the following, perhaps unsatisfactory, possibility. Suppose epistemologists just *define* knowledge as justified true belief—so that these two concepts are functionally equivalent—and suppose that they are extremely rigid in this definition. If the entire epistemology community were to agree upon this definition of knowledge, then the Gettier case would not be a counterexample to a theory of knowledge *as* justified true belief because, in each instance, the individual has justified true belief and *therefore* (by definition) must have knowledge. The Gettier case is successful because of an intuition that is pumped in the vast majority of philosophers—i.e., that there *is not* knowledge in these instances. But, again, this says more about philosophers than it does about knowledge.

Considering again what the trolley problem is supposed to show us, the conclusion to draw from is (obviously) not that it *is* moral to do *X* in such-and-such a scenario and then to do *Y* in the other (or vice-versa). Instead, the conclusion is best represented as a conditional statement; namely,

IF your intuitions are roughly such-and-such,
THEN *X* *explains* why they are so.

Where *X* is filled in with some philosophical analysis.

However, what happens if intuitions diverge? Does this render Foot's analysis false? Or, is the divergent opinion false? Both of these questions are misguided. There is no 'right' and 'wrong' answer in a moral dilemma; there are only intuitions and explanations or theories about the causes of those intuitions. To make the point in the most obvious possible way: a moral dilemma *is a dilemma*; it has no clear solution by design—or rather, it poses a problem that is inherently difficult, by

design. Instead, moral thought experiments serve to perturb the initial conditions of a moral situation until one’s intuitions vary. This can be philosophically useful insofar as we may be able to analyse *which* salient feature of the dilemma caused that variance.

In the case of machine ethics, however, we have seen that moral dilemmas like the trolley problem are used as a *validation proxy* so that ‘if the implementation can resolve a dilemma in a particular manner, then it is deemed to be a successful implementation of ethics in the robot/software agent’ (Nallur, 2020, 2382). This is a category mistake. Moral dilemmas do not tell us what the truth is about whether a particular action is ethical or not; rather, they serve to create new avenues for inquiry. Regardless of one’s metaphilosophical views concerning the ultimate purpose of thought experiments, moral dilemmas have no right answer by design. To suggest that agreement on ethical decision-making in trolley problems is *prior* to moral theorising about AI application presupposes that we have already settled important meta-ethical questions.¹¹

4. SOME META-ETHICAL CONSIDERATIONS

An anti-realist about ethics may, at this point, protest that there are no objective matters of fact about ethics. Therefore, using human data from decisions in moral dilemmas as a benchmark for AI systems is certainly the closest we can get to measuring ethical behaviour—namely, maximising social acceptability. So, the Moral Machine Experiment data is the correct tool for this job.

This is true, but it is also beside the point. Although some authors appear to be sensitive to the targets of their benchmark—i.e., the extent to which, all things considered, a human would accept the decision that an AI system made—it is much more common for there to be a significant conceptual gap between perceived targets and the actual targets of this research. What is problematic here is that researchers often appear to imagine that they are getting at one thing (‘facts’ of ethics) when

¹¹For example, if it were determined that a utility calculus is the ‘correct’ normative theory, then we *could* use moral dilemmas as a validation tool. However, no such determination has been made.

they are really getting at another (sociological facts). It is perceived and therefore presented as though it is the former. This constitutes a derangement of the concept by which, over time, it comes to stand in for the thing itself—this will be a problem as we advance, for all the same reasons that any algorithmic bias is a serious social and philosophical problem.

5. CONCLUSION

Moral dilemmas are used in machine learning to provide circumstances where no ethical option is available to the (artificial) agent making the decision. However, to say that we want an autonomous system to minimise the unethical outcomes under these circumstances presupposes that we already know what the unethical outcomes to be minimised are—i.e., that we have already sorted out the relevant meta-ethical questions.¹² For example, utilitarianism might give a straightforward (and occasionally morally repugnant) answer when deciding between individuals and groups; however, it is less obvious how to calculate what potential future expected utility of a doctor's life will generate that a criminal's life will not. Never mind the fact that it is fallacious to suppose that because most people do reason this way, AI systems ought to reason this way; even if such a calculation is possible, it will always be relative to some frame—increased utility for whom?

Using trolley-style problems in the context of autonomous vehicles as a case study, I have argued that researchers engaged in projects seeking to benchmark ethics are not measuring what they take themselves to be measuring. As we have seen, moral dilemmas are taken to provide something like a ground truth against which an algorithm can be benchmarked. But, this approach to ethical AI systems fails to appreciate the purpose of philosophical thought experiments in the first place. Lack of awareness of this fact sets a dangerous precedent for work in AI

¹²Note that in response to the problem of enabling autonomous systems to distinguish between available choices and to choose the 'least unethical' one, [Dennis et al. \(2016\)](#) suggest that the pressing question to be resolved is 'how can we constrain the unethical actions of autonomous systems but allow them to make justifiably unethical choices under certain circumstances?' But this presupposes that we already know what it means for a decision to be the least unethical. As is common, [Dennis et al. \(2016\)](#) seem to understand 'least unethical' in terms of 'least unacceptable' by the standards of some subset of society.

ethics, because these views get mutually reinforced within the field, leading to a negative feedback loop. The actual target(s) of AI ethics, by dint of being in the realm of moral philosophy, are already highly opaque. The more entrenched the approach of benchmarking ethics using moral dilemmas becomes, as a community-accepted standard, the less clearly individual researchers will see how and why it fails.

This also sets a dangerous precedent when we consider that the majority of AI research is now being done ‘in industry’ (for profit) rather than in academia. Suppose a community-accepted standard for calling, e.g., an autonomous vehicle ‘ethical’ is that it performs well on a set of trolley-style problems, which have been entrenched within the research community as an acceptable benchmark. As noted above, what is actually being measured is how well the machine accords with some set of humans on average, not how ethical the machine actually is—relative to some meta-ethical standard.

This is not to say that moral dilemmas are never inappropriate in the context of AI systems. However, as with any system that uses proxies for measuring alignment of objectives, rather than the objectives themselves, it will be increasingly important that (1) the proxies used are actually representative of the true target, and (2) researchers are aware of what they are actually measuring. Of course, much more work needs to be done in the field of machine ethics to understand the relevant proxies for moral decision-making. Even so, it should be clear that using trolley-style problems (or moral dilemmas more generally) as an elucidatory tool is neither prior to, nor follows from, moral theorising about AI applications.

CONFLICT OF INTEREST

On behalf of all authors, the corresponding author states that there is no conflict of interest.

REFERENCES

Agrawal, Mayank, Joshua C. Peterson, and Thomas L. Griffiths (2020). Scaling up Psychology via Scientific Regret Minimization. *Proceedings of the National Academy of*

- Sciences of the United States of America*, 117(16): 8825–8835.
- Allen, Colin, Wendell Wallach, and Iva Smit (2011). Why Machine Ethics? In Anderson, Michael and Susan Leigh Anderson, editors, *Machine Ethics*, pages 51–61. Cambridge University Press, Cambridge.
- Anderson, Michael and Susan Leigh Anderson (2008). Ethical Healthcare Agents. In Sordo, Margarita, Sachin Vaidya, and Lakhmi C. Jain, editors, *Advances Computational Intelligence Paradigms in Healthcare 3*, volume 107 of *Studies in Computational Intelligence*, pages 233–257. Springer, Berlin, Heidelberg.
- Anderson, Michael, Susan Leigh Anderson, and Chris Armen (2006). MedEthEx: A prototype medical ethics advisor. In *Proceedings of the eighteenth conference on innovative applications of artificial intelligence (IAAI-06)*, Boston, MA. AAAI.
- Arkin, Ronald C. (2008a). Governing Lethal Behavior: Embedding Ethics in a Hybrid Deliberative/Reactive Robot Architecture - Part I: Motivation and Philosophy. In *HRI '08: Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction*, pages 121–128. Association for Computing Machinery.
- Arkin, Ronald C. (2008b). Governing Lethal Behavior: Embedding Ethics in a Hybrid Deliberative/Reactive Robot Architecture - Part II: Formalization for Ethical Control. In *Proceedings of the 2008 Conference on Artificial General Intelligence 2008: Proceedings of the First AGI Conference*, pages 51–62. IOS Press.
- Asaro, Peter (2020). Autonomous Weapons and the Ethics of Artificial Intelligence. In Liao, S. Matthew, editor, *Ethics of Artificial Intelligence*, pages 212–236. Oxford University Press, Oxford.
- Asikainen, Mervi A. and Pekka E. Hirvonen (2014). Thought Experiments in Science and in Science Education. In Matthews, Michael R., editor, *International Handbook of Research in History, Philosophy and Science Teaching*, pages 1235–1256. Springer, Dordrecht.
- Awad, Edmond, Sohan Dsouza, Jean-François Bonnefon, Azim Shariff, and Iyad Rahwan (2020). Crowdsourcing moral machines. *Communications of the ACM*, 63(3): 48–55.
- Awad, Edmond, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan (2018). The Moral Machine Experiment. *Nature*, 563: 59–64.
- Banks, Jaime (2021). Good Robots, Bad Robots: Morally Valenced Behavior Effects on Perceived Mind, Morality, and Trust. *International Journal of Social Robotics*, 13: 2021–2038.
- Bauer, William A. (2020). Virtuous vs. Utilitarian Artificial Moral Agents. *AI and Society*, 35(1): 263–271.
- Baum, Kevin, Holger Hermanns, and Timo Speith (2019). Towards a Framework Combining Machine Ethics and Machine Explainability. In Finkbeiner, Bernd and Samantha Kleinberg, editors, *Third International Workshop on Formal Reasoning about Causation, Responsibility, and Explanations in Science and Technology (CREST 2018)*, volume 286, pages 34–49. Electronic Proceedings in Theoretical Computer Science, EPTCS.
- Bentzen, Martin Mose (2016). The Principle of Double Effect Applied to Ethical Dilemmas of Social Robots. In Seibt, Johanna, Marco Nørskov, and Søren Schack Andersen, editors, *Frontiers in Artificial Intelligence and Applications*, volume 290, pages 268–279. IOS Press, Amsterdam.
- Berreby, Fiona, Gauvain Bourgne, and Jean-Gabriel Ganascia (2015). Modelling Moral Reasoning and Ethical Responsibility with Logic Programming. In Davis, Martin, Ansgar Fehner, Annabelle McIver, and Andrei Voronkov, editors, *LPAR 2015: Logic for Programming, Artificial Intelligence, and Reasoning*, volume 9450 of *Lecture Notes in Computer Science*, pages 532–548. Springer, Berlin, Heidelberg.

- Bhargava, Vikram and Tae Wan Kim (2017). Autonomous vehicles and moral uncertainty. In Lin, Patrick, Keith Abney, and Ryan Jenkins, editors, *Robot Ethics 2.0: From Autonomous Cars to Artificial Intelligence*, pages 5–19. Oxford University Press, Oxford.
- Bjørgen, Edvard P., Simen Madsen, Therese S. Bjørknes, Fredrik V. Heimsæter, Robin Håvik, Morten Linderud, Per-Niklas Longberg, Louise A. Dennis, and Marija Slavkovic (2018). Cake, Death, and Trolleys: Dilemmas as benchmarks of ethical decision-making. In Furman, Jason, Gary Marchant, Huw Price, and Francesca Rossi, editors, *AIES 2018 - Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 23–29. Association for Computing Machinery.
- Blake, P. R., K. McAuliffe, J. Corbit, T. C. Callaghan, O. Barry, A. Bowie, L. Kleutsch, K. L. Kramer, E. Ross, H. Vongsachang, R. Wrangham, and F. Warneken (2015). The ontogeny of fairness in seven societies. *Nature*, 528(7581): 258–261.
- Bojarski, Mariusz, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba (2016). End to end learning for self-driving cars. *arXiv pre-print*, 1604.07316: 1–9. <https://arxiv.org/abs/1604.07316>.
- Bonnefon, J.-F., A. Shariff, and I. Rahwan (2016). The social dilemma of autonomous vehicles. *Science*, 352(6293): 1573–1576.
- Bonnemains, Vincent, Claire Saurel, and Catherine Tessier (2018). Embedded Ethics: Some Technical and Ethical Challenges. *Ethics and Information Technology*, 20: 41–58.
- Bourget, David and David J. Chalmers (2014). What Do Philosophers Believe? *Philosophical Studies*, 170: 465–500.
- Brown, James Robert and Yiftach Fehige (2019). Thought Experiments. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2019 edition.
- Brun, Georg (2018). Thought Experiments in Ethics. In Stuart, Michael T., Yiftach Fehige, and James Robert Brown, editors, *The Routledge Companion to Thought Experiments*, pages 195–210. Routledge, London and New York.
- Bughin, Jacques, Jeongmin Seong, James Manyika, Michael Chui, and Raoul Joshi (2019). Notes from the AI Frontier: Modeling the Impact of AI on the World Economy. *McKinsey Global Institute*.
- Casey, Bryan (2017). Amoral Machines, or: How Roboticians Can Learn to Stop Worrying and Love the Law. *Northwestern University Law Review*, 111(5): 1347–1366.
- Christian, Brian (2020). *The Alignment Problem: Machine Learning and Human Values*. W. W. Norton & Company, New York.
- Ciaramelli, E., M. Muccioli, E. Ladavas, and G. di Pellegrino (2007). Selective deficit in personal moral judgment following damage to ventromedial prefrontal cortex. *Social Cognitive and Affective Neuroscience*, 2(2): 84–92.
- Cointe, Nicolas, Grégory Bonnet, and Olivier Boissier (2017). Jugement éthique dans le processus de décision d’un agent BDI. *Revue d’Intelligence Artificielle*, 31(4): 471–499.
- Conti, Adelaide, Elena Azzalini, Cinzia Amici, Valter Cappellini, Rodolfo Faglia, and Paola Delbon (2017). An Ethical Reflection on the Application of Cyber Technologies in the Field of Healthcare. In Ferraresi, Carlo and Giuseppe Quaglia, editors, *Advances in Service and Industrial Robotics. Proceedings of the 26th International Conference on Robotics in Alpe-Adria-Danube Region, RAAD 2017*, volume 49 of *Mechanisms and Machine Science*, pages 870–876. Springer, Cham.
- Cunneen, Martin, Martin Mullins, Finbarr Murphy, and Seán Gaines (2019). Artificial Driving Intelligence and Moral Agency: Examining the Decision Ontology of Unavoidable Road Traffic Accidents through the Prism of the Trolley Dilemma. *Applied Artificial Intelligence*, 33(3): 267–293.
- Cushman, F., L. Young, and M. Hauser (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. *Psychological Science*, 17(12): 1082–1089.

- Danielson, Peter (2015). Surprising judgments about robot drivers: Experiments on raising expectations and blaming humans. *Etikk i praksis. Nordic Journal of Applied Ethics*, 9(1): 73–86.
- Dennett, Daniel C. (1980). The milk of human intentionality. *The Behavioral and Brain Sciences*, 3: 428–430.
- Dennett, Daniel C. (1991). *Consciousness Explained*. Little, Brown and Co., Boston.
- Dennett, Daniel C. (2013). *Intuition Pumps: and Other Tools for Thinking*. W. W. Norton & Company, New York and London.
- Dennis, Louise, Michael Fisher, Marija Slavkovik, and Matt Webster (2016). Formal Verification of Ethical Choices in Autonomous Systems. *Robotics and Autonomous Systems*, 77: 1–14.
- Döring, Nicola, M. Rohangis Mohseni, and Roberto Walter (2020). Design, Use, and Effects of Sex dolls and Sex Robots: Scoping Review. *Journal of Medical Internet Research*, 22(7): e18551.
- Eichenberg, Christiane, Marwa Khamis, and Lisa Hübner (2019). The Attitudes of Therapists and Physicians on the Use of Sex Robots in Sexual Therapy: Online Survey and Interview Study. *Journal of Medical Internet Research*, 21(8): e13853.
- Etienne, Hubert (2020). When AI Ethics Goes Astray: A Case Study of Autonomous Vehicles. *Social Science Computer Review*. <https://doi.org/10.1177/0894439320906508>.
- Etzioni, Amitai and Oren Etzioni (2017). Incorporating Ethics into Artificial Intelligence. *Journal of Ethics*, 21(4): 403–418.
- Evans, Katherine, Nelson de Moura, Stéphane Chauvier, Raja Chatila, and Ebru Dogan (2020). Ethical Decision Making in Autonomous Vehicles: The AV Ethics Project. *Science and Engineering Ethics*, 26(6): 3285–3312.
- Falbo, Arianna and Travis LaCroix (2021). Est-ce que vous compute? Code-Switching, Cultural Identity, and AI. *arXiv pre-print*, 2112.08256: 1–19. <https://arxiv.org/abs/2112.08256>.
- Foot, Philippa (1967). The Problem of Abortion and the Doctrine of Double Effect. *The Oxford Review*, 5: 5–15.
- Gettier, Edmund L. (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6): 121–123.
- Gordon, John-Stewart (2020). Building Moral Robots: Ethical Pitfalls and Challenges. *Science and Engineering Ethics*, 26(1): 141–157.
- Greene, J., S. Morelli, K. Lowenberg, L. Nystrom, and J. Cohen (2008). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*, 107(3): 1144–1154.
- Greene, Joshua D. (2017). The rat-a-gorical imperative: Moral intuition and the limits of affective learning. *Cognition*, 167: 66–77.
- Greene, J. D., L. E. Nystrom, A. D. Engell, J. M. Darley, and J. D. Cohen (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron*, 44(2): 389–400.
- Greene, J. D., R. B. Sommerville, L. E. Nystrom, J. M. Darley, and J. D. Cohen (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, 293(5537): 2105–2108.
- Grinbaum, Alexei (2018). Chance as a Value for Artificial Intelligence. *Journal of Responsible Innovation*, 5(3): 353–360.
- Harris, John (2020). The Immoral Machine. *Cambridge Quarterly of Healthcare Ethics*, 29(1): 71–79.
- Hauser, M., F. Cushman, L. Young, R. K. Jin, and J. Mikhail (2007). A dissociation between moral judgments and justifications. *Mind and Language*, 22(1): 1–21.
- Headleand, Christopher James, William J. Teahan, and Llyr ap Cenydd (2020). Sexbots: A Case for Artificial Ethical Agents. *Connection Science*, 32(2): 204–221.
- Hellström, Thomas (2013). On the Moral Responsibility of Military Robots. *Ethics and Information Technology*, 15(2): 99–107.

- Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, and Richard McElreath (2001). In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies. *The American Economic Review*, 91(2): 73–78.
- House, Bailey R., Joan B. Silk, Joseph Henrich, H. Clark Barrett, Brooke A. Scelza, Adam H. Boyette, Barry S. Hewlett, Richard McElreath, and Stephen Laurence (2013). Ontogeny of Prosocial Behavior across Diverse Societies. *Proceedings of the National Academy of Sciences of the United States of America*, 110(36): 14586–14591.
- Jobin, Anna, Marcello Ienca, and Effy Vayena (2019). The Global Landscape of AI Ethics Guidelines. *Nature: Machine Intelligence*, 1: 389–399.
- Keeling, Geoff (2018). Legal Necessity, Pareto Efficiency and Justified Killing in Autonomous Vehicle Collisions. *Ethical Theory and Moral Practice*, 21: 413–427.
- Kim, Richard, Max Kleiman-Weiner, Andrés Abeliuk, Edmond Awad, Sohan Dsouza, Joshua B. Tenenbaum, and Iyad Rahwan (2018). A Computational Model of Commonsense Moral Decision Making. In Furman, Jason, Gary Marchant, Huw Price, and Francesca Rossi, editors, *AIES 2018 - Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 197–203. Association for Computing Machinery.
- Koenigs, Michael, Liane Young, Ralph Adolphs, Daniel Tranel, Fiery Cushman, Marc Hauser, and Antonio Damasio (2007). Damage to the pre-frontal cortex increases utilitarian moral judgements. *Nature*, 446(7138): 908–911.
- Krishnan, Armin (2009). *Killer Robots: Legality and Ethicality of Autonomous Weapons*. Ashgate, Surrey.
- Krylov, Nikolay N., Yevgeniya L. Panova, and Aftandil V. Alekberzade (2019). Artificial Morality for Artificial Intelligence. *History of Medicine*, 6(4): 191–199.
- Kuefler, Alex, Jeremy Morton, Tim Wheeler, and Mykel Kochenderfer (2017). Imitating Driver Behavior with Generative Adversarial Networks. *IEEE Intelligent Vehicles Symposium*, IV: 204–211.
- Kuhn, Thomas S. (1977). A Function for Thought Experiments. In *The Essential Tension: Selected Studies in Scientific Tradition and Change*, pages 240–265. University of Chicago Press, Chicago.
- LaCroix, Travis and Aydin Mohseni (2020). The Tragedy of the AI Commons. *arXiv pre-print*, 2006.05203: 1–40. <https://arxiv.org/abs/2006.05203>.
- Lim, Hazel Si Min and Araz Taeihagh (2019). Algorithmic Decision-Making in AVs: Understanding Ethical and Technical Concerns for Smart Cities. *Sustainability*, 11(20): 5791.
- Lin, Patrick (2015). Why Ethics Matters for Autonomous Cars. In Maurer, M., J. Gerdes, B. Lenz, and H. Winner, editors, *Autonomes Fahren*, pages 69–85. Springer Vieweg, Berlin and Heidelberg.
- Lindner, Felix, Martin Mose Bentzen, and Bernhard Nebel (2017). The HERA approach to morally competent robots. In *IEEE International Conference on Intelligent Robots and Systems*, volume 2017-September, pages 6991–6997. IEEE.
- Lindner, Felix, Robert Mattmüller, and Bernhard Nebel (2020). Evaluation of the Moral Permissibility of Action Plans. *Artificial Intelligence*, 287: 103350.
- Lourie, Nicholas, Ronan Le Bras, and Yejin Choi (2020). SCRUPLES: A Corpus of Community Ethical Judgments on 32,000 Real-life Anecdotes. *arXiv pre-print*, 2008.09094: 1–16. <https://arxiv.org/abs/2008.09094>.
- Luccioni, Alexandra and Yoshua Bengio (2019). On the Morality of Artificial Intelligence. *arXiv pre-print*, 1912.11945: 1–12. <https://arxiv.org/abs/1912.11945>.
- Malle, Bertram F., Matthias Scheutz, Thomas Arnold, John Voiklis, and Corey Cusimano (2015). Sacrifice One for the Good of Many?: People Apply Different Moral Norms to Human and Robot Agents. In *HRI '15: Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 117–124. Association for Computing Machinery.
- Mayo-Wilson, Conor and Kevin J.S. Zollman (2020). The Computational Philosophy: Simulation as a Core Philosophical Method. *PhilSci Archive*, 18100: 1–32. <http://>

- philsci-archive.pitt.edu/18100/.
- Misselhorn, Catrin (2018). Artificial Morality. Concepts, Issues and Challenges. *Society*, 55(2): 161–169.
- Moore, A., B. Clark, and M. Kane (2008). Who shalt not kill? Individual differences in working memory capacity, executive control, and moral judgment. *Psychological Science*, 19(6): 549–557.
- Nallur, Vivek (2020). Landscape of Machine Implemented Ethics. *Science and Engineering Ethics*, 26(5): 2381–2399.
- Navarrete, C. David, Melissa M. McDonald, Michael L. Mott, and Benjamin Asher (2012). Virtual Morality: Emotion and Action in a Simulated Three-Dimensional ‘Trolley Problem’. *Emotion*, 12(2): 364–370.
- Nichols, S. and R. Mallon (2005). Moral dilemmas and moral rules. *Cognition*, 100(3): 530–542.
- Noothigattu, Ritesh, Snehal Kumar (Neil) S. Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, and Ariel D. Procaccia (2018). A Voting-Based System for Ethical Decision Making. In *The Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 1587–1594. Association for the Advancement of Artificial Intelligence.
- Olson, Randal S., William La Cava, Patryk Orzechowski, Ryan J. Urbanowicz, and Jason H. Moore (2017). PMLB: a large benchmark suite for machine learning evaluation and comparison. *BioData Mining*, 10(36): 1–13.
- Pardo, Antonio Miguel Seoane (2018). Computational Thinking between Philosophy and STEM - Programming Decision Making Applied to the Behavior of ‘Moral Machines’ in Ethical Values Classroom. *Revista Iberoamericana de Tecnologías del Aprendizaje*, 13(1): 20–29.
- Pereira, Luís Moniz and Ari Saptawijaya (2011). Modeling Morality with Prospective Logic. In Anderson, Michael and Susan Leigh Anderson, editors, *Machine Ethics*, pages 398–421. Cambridge University Press, Cambridge.
- Pereira, Luís Moniz and Ari Saptawijaya (2015). Bridging Two Realms of Machine Ethics. In White, Jeffrey and Rick Searle, editors, *Rethinking Machine Ethics in the Age of Ubiquitous Technology*, pages 197–224. Information Science Reference, Hershey, PA.
- Raji, Inioluwa Deborah, Emily M. Bender, Amandalynne Paullada, Emily Denton, and Alex Hanna (2021). AI and the Everything in the Whole Wide World Benchmark. *arXiv pre-print*, 2111.15366: 1–20. <https://arxiv.org/abs/2111.15366>.
- Russell, Stuart (2019). *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking, New York.
- Sans, Alger and David Casacuberta (2019). Remarks on the Possibility of Ethical Reasoning in an Artificial Intelligence System by Means of Abductive Models. In Nepomuceno-Fernández, Ángel, Lorenzo Magnani, Francisco J. Salguero-Lamillar, Cristina Barés-Gómez, and Matthieu Fontaine, editors, *MBR 2018: Model-Based Reasoning in Science and Technology*, volume 49 of *Studies in Applied Philosophy, Epistemology and Rational Ethics*, pages 318–333. Springer, Cham.
- Santoni de Sio, F. (2017). Killing by Autonomous Vehicles and the Legal Doctrine of Necessity. *Ethical Theory and Moral Practice*, 20(2): 411–429.
- Saptawijaya, Ari and Luís Moniz Pereira (2015). Logic Programming Applied to Machine Ethics. In Pereira, Francisco, Penousal Machado, Ernesto Costa, and Amílcar Cardoso, editors, *EPIA 2015: Progress in Artificial Intelligence*, volume 9273 of *Lecture Notes in Computer Science*, pages 414–422. Springer, Cham.
- Saptawijaya, Ari and Luís Moniz Pereira (2016). Logic programming for modeling morality. *Logic Journal of the IGPL*, 24(4): 510–525.
- Schaich Borg, J., C. Hynes, J. J. Van Horn, S. Grafton, and W. Sinnott-Armstrong (2006). Consequences, action, and intention as factors in moral judgments: An fMRI investigation. *Journal of Cognitive Neuroscience*, 18(5): 803–817.

- Sharkey, Amanda and Noel Sharkey (2012). Granny and the robots: ethical issues in robot care for the elderly. *Ethics and Information Technology*, 14(1): 27–40.
- Sommaggio, Paolo and Samuela Marchiori (2018). Break The Chains: A New Way To Consider Machine’s Moral Problems. *BioLaw Journal*, 2018(3): 241–257.
- Sütfeld, Leon R., Richard Gast, Peter König, and Gordon Pipa (2017). Using Virtual Reality to Assess Ethical Decisions in Road Traffic Scenarios: Applicability of Value-of-Life-Based Models and Influences of Time Pressure. *Frontiers in Behavioral Neuroscience*, 11: 122.
- Szczepański, Marcin (2019). Economic Impacts of Artificial Intelligence (AI). *European Parliamentary Research Service*, PE 637.967: 1–8.
- Thomson, Judith Jarvis (1976). Killing, Letting Die, and the Trolley Problem. *The Monist*, 59: 204–217.
- Tonkens, Ryan (2012). The Case Against Robotic Warfare: A Response to Arkin. *Journal of Military Ethics*, 11(2): 149–168.
- Unger, Peter (1996). *Living and Letting Die*. Oxford University Press, Oxford.
- Vincent, James (2018). Global Preferences for Who to Save in Self-driving Car Crashes Revealed: Congratulations to young people, large groups of people, and people who aren’t animals. *The Verge*, 24 Oct 2018. <https://www.theverge.com/2018/10/24/18013392/self-driving-car-ethics-dilemma-mit-study-moral-machine-results>.
- Wakabayashi, Daisuke (2018). Self-Driving Uber Car Kills Pedestrian in Arizona, Where Robots Roam. *The New York Times*, 19 Mar 2018. <https://www.nytimes.com/2018/03/19/technology/uber-driverless-fatality.html>.
- Waldmann, M. R. and J. H. Dieterich (2007). Throwing a bomb on a person versus throwing a person on a bomb: Intervention myopia in moral intuitions. *Psychological Science*, 18(3): 247–253.
- Wallach, Wendell and Colin Allen (2009). *Moral Machines: Teaching Robots Right from Wrong*. Oxford University Press, Oxford.
- Welsh, Sean (2017). *Ethics and Security Automata: Policy and Technical Challenges of the Robotic Use of Force*. Routledge, London and New York.
- Wintersberger, Philipp, Anna-Katharina Frison, Andreas Riener, and Shailie Thakkar (2017). Do moral robots always fail? Investigating human attitudes towards ethical decisions of automated systems. In *26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 1438–1444. IEEE.
- Wright, Ava Thomas (2019). Rightful Machines and Dilemmas. In Conitzer, Vincent, Gillian Hadfield, and Shannon Vallor, editors, *AIES ’19 - Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 3–4. Association for Computing Machinery.
- Yang, Shun, Wenshuo Wang, Chang Liu, Weiwen Deng, and J. Karl Hedrick (2017). Feature Analysis and Selection for Training an End-to-end Autonomous Vehicle Controller Using Deep Learning Approach. In *IEEE Intelligent Vehicles Symposium*, volume IV, pages 1033–1038. IEEE.