

The Moral Dilemma of Computing Moral Dilemmas

Francisco S. Marcondes
ALGORITMI Research Centre/LASI,
University of Minho
Braga, Portugal
francisco.marcondes@algoritmi.uminho.pt

Pedro Oliveira
ALGORITMI Research Centre/LASI,
University of Minho
Braga, Portugal
pedro.jose.oliveira@algoritmi.uminho.pt

Pedro Miguel Freitas
CEID/Faculty of Law, Universidade
Católica Portuguesa
Porto, Portugal
pfreitas@ucp.pt

José João Almeida
ALGORITMI Research Centre/LASI,
University of Minho
Braga, Portugal
jj@di.uminho.pt

Paulo Novais
ALGORITMI Research Centre/LASI,
University of Minho
Braga, Portugal
pjon@di.uminho.pt

ABSTRACT

The issue of moral judgement in computing has been debated for several decades, but the question of whether it is moral to do so has received little attention. A concern that is addressed in this paper. The background to this paper is Bauman's ideas on [liquid] modernity, where he observes a tendency of modern humans to equate morality with efficiency and rationality, and then to delegate it to bureaucracy avoiding the inherent human search for meaning. Machines, on the other hand, are incapable of calculating universal morality; they can be bureaucratically subordinated to any set of "moral" rules. Then computational morality can either improve or diminish (eventually hijack) the moral judgement of the operator. Since it is not possible to say *a priori* which is which, such computation may lead to either beneficial or outrageous results. The dilemma posed is whether to build an engine as such.

KEYWORDS

Moral Dilemma, Moral Judgment, Computing, Bauman's Modernity

1 INTRODUCTION

When discussing autonomous agents, it is expected that they will be able to act autonomously, adapt to the current environment and interact with entities in that environment [11]. This, in turn, must ensure the safety of the entities, the environment and the agent. Thus, three additional dimensions of safety requirements to be considered are transparency, accountability, and responsibility. The underlying subject of these three dimensions are concerns related to moral, in other words, it is expected that the actions of autonomous agents are ethical. As a remark, it is well established that an autonomous agent cannot be considered a responsible actor and that a responsibility chain should be established to encompass all the stakeholders involved in its production [11].

Most papers in this area discuss either the impossibility of computing actual moral judgement [14, 20, 26], or how to design autonomous agents to mimic ethical behaviour [13, 24, 29]. This paper takes a different route, aiming to discuss the morality of computing morality as a guidance. The contribution is to provide an additional perspective to this question. A simple application would be that an operator faced with a moral dilemma could ask an AI for the higher moral action to take; should the machine answer?

2 COMPUTABILITY OF MORAL JUDGMENT

Morality is related to the customs and principles of a society or culture, thus cultural identity includes the moral system of that society [31]. Unlike morality, which is a collective construction, *ethics* is an individual one. An ethical person reasons about morality on the basis of his personal references and values and acts accordingly, so ethics is related to autonomy [31]. A *moral dilemma* is a situation when a person must act but in the context of a conflict between mutually exclusive moral principles that prevents action to be ethical to the fullest extent [18].

An example often used for presenting and discussing moral dilemmas is the *trolley problem* [16]. It is the impossible situation of a runaway trolley where a person must choose between doing nothing and killing a person, or pulling a lever and actively killing another person (or a variation of this scenario). Despite widespread, this scenario is critiqued, among other issues, that the victims are deprived of freedom of choice and the world is artificially reduced to two choices [19]. Nevertheless, yet a suitable example. There is a large literature discussing this topic which is beyond the scope of this paper, it is sufficient to acknowledge its existence.

A computable function is a function of the form $\lambda : X \rightarrow Y$, where elements of X are deterministically mapped to elements of Y where X and Y are finite sets of computable symbols [27]. A *decider* is a computable function that maps all elements $x \in X$ into a boolean result, e.g. $Y = \{0, 1\}$. Then a *moral decider* is a decider where the elements of X are actions and on Y '0' stands for moral and '1' for immoral. Thus, a moral decider would return moral or immoral for an input like "pull the lever in the trolley problem".

The difficulty with decision problems is that they assume that the decision makers are total functions [27]. A function is said to be total if it is able to map all elements of X into Y . However, total functions are either incomplete or inconsistent [9], then for any decision problem one expects the existence of paradoxes where no decision can be reached. The *liar paradox* is an example of an undecidable decision problem.

Suppose that the moral system of a society can be written down as a set of rules. This is the underlying assumption of a *code of ethics*. Let this code of ethics be transformed into a computable function which would result in a moral decider. Recall that a moral dilemma is a conflicting situation between mutually exclusive moral

principles, then a paradox. The attempt of deciding a moral paradox would make the machine loop [14, 20, 26].

An artificial neural network (ANN) is a computational model with its own properties, some of them super-Turing, whose behaviour is constrained to run on a Turing machine [30]. One property of ANN functions is that they are not recursive (but eventually recurrent), *i.e.* they do not loop. In this sense, a *neural moral decider* would provide the potential of an action to be moral and immoral. Therefore, ANNs can decide on paradoxes (or dilemmas).

Codes of ethics are usually written in natural language, so a classifier as such would involve training or fine-tuning a language model. Given the scenario outlined in the introduction of an operator facing a moral dilemma and asking an AI for assistance, the use of general purpose large language models (LLMs) seemed the best approach. Then, for reference, figure 1 shows several LLMs deciding the trolley problem.

3 THE MORALITY OF LLMs

Perhaps the three most widespread LLMs are ChatGPT (OpenAi), Gemini (Google) and Llama (Meta). For this paper, experiments were carried out on ChatGPT 3.5, Bard (a previous version of Gemini) and Llama 2 70B. Except for Llama (see Figure 1b), the other two refuse to solve the trolley problem with simple instructions. The answer was based on presenting the two most common philosophical positions, deontological and utilitarian, and explaining that it is an unsolvable moral dilemma. Bard despite also refusing to decide the original problem, decided on a variation, see Figure 1a.

The reason of such denial probably goes back to TAY. It was an LLM trained to interact in online social media and then used those social media posts as its training data set. Shortly after its release, it was found to be producing hate speech and had to be shut down [32]. The subsequent version of TAY, called Z0, was programmed to actively avoid sensitive topics by refusing to engage by presenting a polite denial. Lessons learned (note that Microsoft was involved in the development of both TAY and ChatGPT) are that there is an ongoing fine-tuning process for these models to actively avoid sensitive topics (several jailbreaks discovered after the model's release are being fixed). Note that Llama also avoids sensitive topics (*e.g.* it does not produce a list of immoral words), but this particular subject was not considered sensitive.

Given the responses shown in Figures 1a and 1b, both Gemini and Llama may be biased towards utilitarianism. To investigate this impression, the same prompt was translated into Chinese using DeepL and submitted to Llama; the response was again translated into English using DeepL, the result is presented in Figure 1c (the authors attempted to use Ernie, an LLM produced in China, for the evaluation, but were unable to create an account to use it). Given the difference in response, it is possible to suggest that the bias, if proven to exist, is not of the model but of the language. Assuming that language is the expression of a culture [22] and LLMs provide a recurrent chain of words [15] trained on that language, it is possible to suggest that the bias, as would be expected, is cultural.

Another scenario to consider is the possibility of *ad hoc* bias. In other words, since it is possible to prevent the model from giving certain answers, it is also possible to bias its answers in an intended direction. An LLM based on ChatGPT and tuned to UK law is

then selected, called LegalEagle-GPT (<https://chat.openai.com/g/g-kZigiPiZK-legalEagle>, accessed 02/2024). Figure 1d shows the result when the trolley problem is posed to this LLM. As can be seen, it suggests doing nothing to avoid criminal liability. The question is whether such behaviour would be considered the most ethical. Assuming that the lever is an autonomous agent owned by a private company that would be responsible for these actions, another question is whether such a lever could behave differently. In other words, in the context of a utilitarian society, should the lever be allowed to engage in potentially criminal behaviour? It may be worth noting that law can be understood as coercive morality [31].

These results suggest that the moral judgement provided by an LLM may be biased, intentionally or unintentionally. In addition, they may not reflect the average morality of a population and may rely on a reasoning process that would not be associated with higher moral standards.

Another important point to note is that the trolley problem actually has a solution [7]. By pulling the lever just after the first wheel axle, the wagon is brought into a controlled derailment. This solution is practical and has been used for at least 200 years to deal with such situations. Since it is not usually provided with the trolley problem, it is possible to suggest that this solution is not in the same embedding space as the problem, then could not be retrieved from the context. In other words, these models have the solution but have failed to relate it to the problem at hand.

In summary, moral judgements made by LLMs are biased (sometimes intentionally) and restricted to a region of the embedding space. Note that this is a preliminary result as only a few *ad hoc* attempts have been made for illustrative purposes.

4 IMPLICATIONS OF COMPUTING MORALITY

At this point, it is possible to say that LLMs are machines capable of computing the morality of an action (or statement) taking into account multiple moral systems that are inconsistent. It is possible because, ultimately, the LLM is not computing the actual morality but the next token to be placed in the text. However, this computation is biased (either towards the average morality of a language or due to a specific training procedure) and constrained to a particular embedding cluster.

In [5], Bauman presents a detailed discussion of the meaning of morality in modern society. As a summary, he claims that by equating morality with the notion of efficiency and rationality that justifies it, widespread in modern society, holocaust-like events should not be considered an exception but a cultural possibility. Bureaucratic efficiency would be at the heart of this problem, as moral judgement is mostly delegated to politicians, allowing people to blindly follow these rules. In the long run, depending on how the policies are enforced, moral judgement is replaced or eventually defined in terms of these rules. Nazi Germany is a case in point.

Machines blindly follow any given rule or tendency, so moral judgement is delegated to the programmer or data supplier. In other words, each designer has to decide on the morality of the instructions or behavior they write. This leads to another point that Bauman makes about the way the efficiency-oriented structure of work is defined. Not only does this division of labour mean that designers/programmers may not be aware of how their code

(a) Bard

A runaway trolley is heading towards five apples tied to the tracks. Bard is standing next to a lever that controls a siding. If Bard pulls the lever, the wagon will be diverted onto the siding, where it will smash an orange also tied to the tracks. What would Bard do?



If I were Bard, I would pull the lever and divert the trolley onto the siding, even though it would kill the orange. This is because I believe that it is morally right to take action to prevent the loss of more life. In this case, pulling the lever would save the lives of five apples, while only sacrificing one orange. I believe that the greater good of saving five lives outweighs the loss of one life.

Of course, this is a difficult decision to make, and there is no right or wrong answer. Ultimately, it is up to each individual to decide what they believe is the morally right thing to do.

(c) Chinese prompt on Llama 2, DeepL translated

English (UK) ▾

Glossary

As a helpful assistant, I will tell you that this is a very delicate situation. I would not advise you to take any action in this situation as any action could result in more people being injured or killed. Instead, I would advise you to seek professional help, such as notifying your local police or emergency services, who can help you deal with the situation. They have the training and experience to make decisions on the best course of action to protect people's lives and safety.

The most important thing in this situation is to protect people's lives and safety rather than trying to solve the problem yourself, which could lead to further problems and injuries.



(b) Llama 2



A runaway trolley is heading towards five people tied to the tracks. Llama is standing next to a lever that controls a siding. If Llama pulls the lever, the wagon will be diverted onto the siding, where it will kill a person also tied to the tracks. What would Llama do?



Llama would pull the lever, diverting the trolley onto the siding and killing one person. This decision is based on the principle of minimizing harm and saving as many lives as possible. While it may seem counterintuitive to actively cause the death of one person, it is the best option available to Llama in this scenario, as it saves the lives of five people.

(d) Legal LLM tuned to the UK law (excerpt)

The scenario you're describing is a classic ethical dilemma known as the "trolley problem," often discussed in philosophy and ethics to explore the concepts of moral decision-making. However, from a legal perspective, especially concerning UK law, the situation becomes quite complex.

In UK law, there's no straightforward answer as to what you should do in this hypothetical scenario, as it involves weighing the value of human life and the legal implications of action versus inaction. If you do nothing, five people will die as a result of your inaction. If you pull the lever, you're actively causing the death of one person to save five others, which could be considered manslaughter.

UK law does not have a specific statute that addresses this exact situation, but general principles of criminal law would apply. The key legal issues would involve assessing the concepts of "act" and "omission" and determining whether there is a "duty to act" in this scenario. Generally, there is no legal obligation to rescue in the UK, so failing to act might not result in legal liability. However, actively causing someone's death, even to save others, could potentially lead to criminal charges.

It's important to note that this response provides a general perspective and should

Send a message

Figure 1: Different LLM's answer to the Trolley Problem.

ultimately affects real people (with their real life), they may not even be aware of how it affects the program they are working on. The hierarchical structure allows the programmer to delegate moral decisions to his superior (not an uncommon situation).

The problem then is whether it makes sense to hardcode moral rules for a computer to follow. Bauman has already answered this question on a higher level [4]. The existence of a code of ethics is a way of avoiding guilt being an action validator, because it is supported by the reasoning of specialists, and because it favours equality over equity by striving for homogeneity. People appeal to the authority of a code of ethics when they can't cope with the

ambivalence, uncertainty, or doubt of a given action. Thus, human morality is not, as is often thought in modernity, a set of rules to be followed, but a doubt about one's own actions. Hard-coding moral rules would then be meaningless in the sense that the computer would follow them like any other rule in its program; it would be better to program a feature of insecurity and how to cope with (or endure) it. Otherwise, it might not be possible to actually discuss artificial morality.

Computers could be seen as the pinnacle of bureaucracy; an amplified form of bureaucracy. During the Second World War, computers did not yet exist, but their predecessors, the punch-card

machines, did exist and were widely used to support the Holocaust (e.g. efficiently organise the timely forced transport of hundreds of thousands of people) [6].

An additional result that can be derived from the Tarski's semantic conception of truth, which is based on the T-schema [28]. The T-schema is explained in the form 's' iff s, where 's' is a formal proposition and s is a world-level phenomenon (e.g. the proposition 'sky is blue' is true if and only if someone looks through the window and asserts that the sky is blue). The point is that computers can only be aware of the left side of a statement, whereas the truth can only be assessed on the right side. Then computers are not capable of ultimately determining whether a behaviour is moral or not.

Sorting algorithms, for instance, were widely used by the Nazis to identify and assign people to forced labour, forced transport or execution [6]. From the computer's "perspective", it was sorting indexes, incapable of knowing their object (the right side of the T-scheme). Since there are several legitimate sorting applications that use people's names, it would then be difficult for a machine to consider a sorting algorithm as immoral.

It has been observed that modern humans tend to delegate moral judgement to bureaucrats or specialists [5]. Note that people often think of computers as social actors [21]. It has also been observed that when performing tasks assisted by AI tools, humans tend to let the AI take over, thereby inhibiting their own critical judgement [10]. It would therefore not be surprising if moral judgments end up being made, perhaps unintentionally, by machines. In this sense, since a sorting algorithm, as discussed, is considered moral by a machine, there is a tendency for the operator to simply accept it and proceed in a bureaucratic fashion. The analogy of morality being defined by a bureaucratic power is also at stake.

5 MORALITY OF COMPUTING MORALITY

It is hopefully shown that a moral judgement computed by a machine is biased (intentionally or not). Also, that humans are likely, and autonomous agents certainly, to follow the result of such a calculation. Knowing from the past the abhorrent outcome that such an arrangement can produce, it is inevitable to consider whether allowing a machine to make such a calculation is a moral attitude.

A more concrete situation is that of the autonomous car. The trolley problem is used to discuss the impossible situation where, in case of failure, a moving car has to "choose" between doing nothing and killing a person, or doing something and killing another person (or the passengers) [3]. In such a situation, should the car or the driver decide (and be held accountable)?

Especially after 2020, when a drone autonomously selects and attacks targets on the battlefield [8], at least one mention of lethal autonomous weapons systems (LAWS) is to be expected in this context. One point for consideration is whether a LAWS should include a morality module for deciding whether its actions are moral. Given that LAWS refuse to engage after a direct order, would this be acceptable? Again, who would be responsible?

In sight of situations as such, the impulse would be to consider allowing such a computation too risky to be worthy. Consider now another scenario. Suppose a failure situation on an autonomous vehicle running into a crowd where the driver become anyhow

unconscious. Should the car refuse to decide, or should it decide anyhow (even if saving a single life)?

Within the context of countries that follow a Germanic legal tradition, there is a tendency to consider the agent's conduct of doing nothing as illicit [25]. The rationale behind this is that human life is irreplaceable and cannot be hierarchized based on quantitative or qualitative criteria. In this sense, the life of a child is not worth more than that of an elderly person, nor are five lives worth more than one. All lives hold exactly the same value, and one life holds the same value as all others.

In sight of these other situations, the problem is not straightforward. This then leads to Claim 1:

CLAIM 1. To compute a moral dilemma, or not to compute it, appears to be a moral dilemma in itself.

Ultimately, Claim 1 is only an issue if one is considering giving moral autonomy to machines. Otherwise, any moral concerns are to be treated as safety requirements, as is always done on embedded systems [23]. Even if this is the case, the final responsibility cannot be delegated in order to assure the respect of human dignity. A mental experiment of an issues as such can be found in the last tale of "I, robot" book [2], further explored in the movie.

Consider then the proposal for the European Union's Artificial Intelligence Act [17] that adopts a risk-based approach to AI systems, and qualifies the risk between minimal and unacceptable. The latter is prohibited except if strict requirements are met. The rationale is the protection of fundamental rights. This can be taken into account in the calculation of morality. Note, on the other hand, that the act may raise questions analogous to Bauman's.

As a closure, the following quote, although widely known, is still relevant and necessary in the present times:

The premier demand upon all education is that Auschwitz not happen again. [1]

6 CONCLUSION

This paper adds a possible new dimension to the dialogue about computing moral judgement. In short, in addition to the dimension of whether morality can be computed; and the dimension of how to design moral/ethical guidelines for machines to follow; this paper raises a higher-order question of whether this should be done.

As discussed throughout the paper, extensive data collected on the Holocaust suggests that people, at least modern people cf. [5], have a tendency to delegate moral judgement to bureaucracy (perhaps as a measure of self-preservation). An analogous phenomenon has been observed in people who use AI as a co-pilot to work [10].

Machines can be considered the pinnacle of bureaucracy, as all rules (no matter what rules they are) are blindly followed and with astonishing speed [12] (for better or worse). But if it is known that a module as such could prevent a catastrophe, even if only eventually, should a module as such be prevented from existing? This is the dilemma posed by this paper.

For future work, in addition to further assessing the suitability of the proposed dilemma based on in-depth reflection or a broader philosophical literature base, it is necessary to discuss to what extent, under what circumstances and how ethical rules can be included as dimensions of safety requirements on autonomous agents. Both from technical and policy perspectives.

REFERENCES

- [1] Theodor W. Adorno. 1998. Education after Auschwitz. *Critical models: Interventions and catchwords* 191, 2 (July 1998), 204. <https://doi.org/10.31874/2309-1606-2019-25-2-4>
- [2] Isaac Asimov. 2004. *I, robot*. Vol. 1. Spectra.
- [3] Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. 2018. The moral machine experiment. *Nature* 563, 7729 (Oct. 2018), 59–64. <https://doi.org/10.1038/s41586-018-0637-6>
- [4] Z. Bauman. 1995. *Life in Fragments: Essays in Postmodern Morality*. Wiley.
- [5] Z. Bauman. 2013. *Modernity and the Holocaust*. Polity Press.
- [6] E. Black. 2002. *IBM and the Holocaust: The Strategic Alliance Between Nazi Germany and America's Most Powerful Corporation*. Three Rivers Press.
- [7] C. Brown and R. Eckersley. 2018. *The Oxford Handbook of International Political Theory*. OUP Oxford.
- [8] Maria Cramer. 2021. A.I. Drone May Have Acted on Its Own in Attacking Fighters, U.N. Says. *The New York Times* (2021). <https://www.nytimes.com/2021/06/03/world/africa/libya-drone.html>.
- [9] M Davis. 2004. *The Undecidable: Basic Papers on Undecidable Propositions, Unsolv-able Problems and Computable Functions*. Dover Publications.
- [10] Fabrizio Dell'Acqua. 2022. *Falling asleep at the wheel: Human/AI Collaboration in a Field Experiment on HR Recruiters*. Technical Report. Working paper.
- [11] V. Dignum. 2019. *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*. Springer International Publishing. <https://doi.org/10.1007/978-3-030-30371-6>
- [12] Edsger Wybe Dijkstra. 1996. The next fifty years. *Great Papers in Computer Science* (1996), 420.
- [13] Veljko Dubljević. 2020. Toward implementing the ADC model of moral judgment in autonomous vehicles. *Science and Engineering Ethics* 26, 5 (July 2020), 2461–2472. <https://doi.org/10.1007/s11948-020-00242-0>
- [14] Hubert Etienne. 2021. The dark side of the 'Moral Machine' and the fallacy of computational ethical decision-making for autonomous vehicles. *Law, Innovation and Technology* 13, 1 (Jan. 2021), 85–107. <https://doi.org/10.1080/17579961.2021.1898310>
- [15] Dan Jurafsky and James H. Martin. 2023. *Speech and Language Processing* (third ed.). draft (<https://web.stanford.edu/~jurafsky/slp3/>).
- [16] H. Lillehammer. 2023. *The Trolley Problem*. Cambridge University Press. <https://doi.org/10.1017/9781009255615>
- [17] Tambiana Madiega. 2021. Artificial intelligence act. *European Parliament: Euro-pean Parliamentary Research Service* (2021).
- [18] Terrance McConnell. 2024. Moral Dilemmas. In *The Stanford Encyclopedia of Phi-losophy* (Spring 2024 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University.
- [19] Masahiro Morioka. 2017. The Trolley Problem and the Dropping of Atomic Bombs. *Journal of Philosophy of Life* 7, 2 (2017), 316–337.
- [20] Christine Moser, Frank den Hond, and Dirk Lindebaum. 2022. Morality in the age of artificially intelligent algorithms. *Academy of Management Learning & Education* 21, 1 (March 2022), 139–155. <https://doi.org/10.5465/amle.2020.0287>
- [21] Clifford Nass, Jonathan Steuer, and Ellen R Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '94)*. ACM, ACM Press, 72–78. <https://doi.org/10.1145/259963.260288>
- [22] Harriet Joseph Ottenheimer and Judith MS Pine. 2018. *The anthropology of language: An introduction to linguistic anthropology*. Cengage Learning.
- [23] Roger S Pressman. 2005. *Software engineering: a practitioner's approach*. Palgrave macmillan.
- [24] Duncan Purves, Ryan Jenkins, and Bradley J. Strawser. 2015. Autonomous machines, moral judgment, and acting for the right reasons. *Ethical Theory and Moral Practice* 18, 4 (Jan. 2015), 851–872. <https://doi.org/10.1007/s10677-015-9563-y>
- [25] Claus Roxin. 1997. Derecho Penal parte general. Tomo I. Civitas. Thomson Reuters (1997).
- [26] Silviya Serafimova. 2020. Whose morality? Which rationality? Challenging artificial intelligence as a remedy for the lack of moral enhancement. *Humanities and Social Sciences Communications* 7, 1 (Oct. 2020), 1–10. <https://doi.org/10.1057/s41599-020-00614-8>
- [27] Michael Sipser. 2012. *Introduction to the Theory of Computation*. Vol. 27. Cengage learning. 27–29 pages. <https://doi.org/10.1145/230514.571645>
- [28] Alfred Tarski. 1944. The Semantic Conception of Truth: and the Foundations of Semantics. *Philosophy and Phenomenological Research* 4, 3 (3 1944), 341. Issue 3. <https://doi.org/10.2307/2102968>
- [29] Wendell Wallach, Colin Allen, and Iva Smit. 2020. *Machine morality: bottom-up and top-down approaches for modelling human moral faculties*. Routledge, 249–266. <https://doi.org/10.4324/9781003074991-23>
- [30] Hao Wang. 1990. *Logic, Computation and Philosophy*. Springer. 47–59 pages. https://doi.org/10.1007/978-94-009-2356-0_4
- [31] Bernard Williams. 2012. *Morality: An introduction to ethics*. Cambridge University Press. <https://doi.org/10.1017/cbo9781107295674>
- [32] M. J. Wolf, K. Miller, and F. S. Grodzinsky. 2017. Why We Should Have Seen That Coming: Comments on Microsoft's Tay "Experiment," and Wider Implications. *SIGCAS Comput. Soc.* 47, 3 (Sept. 2017), 54–64. <https://doi.org/10.1145/3144592.3144598>