Contents

Preface

"Today we tend to go on for years, with tremendous investments to find that the system, which was not well understood to start with, does not work as anticipated. We build systems like the Wright brothers built airplanes – build the whole thing, push it off the cliff, let it crash, and start over again."
(Graham, in *Software Engineering and Society*, (Naur, 1968).)

1 Methodological Background

The purpose of this book is to engage with historical and philosophical issues underpinning, what we identify here, as 'programming systems', viz. large systems that have been programmed in order to control some process or set of processes.

In a recent paper published in the *Communications of the ACM*, we read the following assessment of the state of modern computing systems, (Neumann, 2017, p. 3):

> *"Unfortunately, the trends for the future seem relatively bleak. Computer system trustworthiness and the implications of its absence are increasingly being questioned. Semi- and fully autonomous systems, the seemingly imminent Internet of Things, and artificial intelligence are providing further examples in which increasing complexity leads to obscure and unexplainable system behavior. The concept of trustworthiness seems to becoming supplanted with people falsely placing their trust in systems and people that are simply not trustworthy – without any strong cases being made for safety, security, or indeed assurance that might otherwise be found in regulated critical industries".*

Trustworthiness of large systems is just one of a growing number of serious problems related to computing, with the potential to affect millions of lives.[1] This is due not just to properties of the systems themselves, but also to their use, design and

[1] This is a long standing issue in computing, touching on several areas. One of the early and most broad views on computing, risk and trust can be found in MacKenzie (2004). Recently, the area of computational trust has grown sensibly in its impact and applications, from software packages distribution systems to vehicular networks, see e.g. Primiero, Boender (2017); Primiero et al. (2017)

development by humans. On the one hand, these systems are ubiquitous, both in terms of usage and impact: almost everyone in large part of the developed world interacts constantly with a computing device; also, some of these systems have progressively evolved into cyber-physical entities, capable of acting upon and being affected by the external environment. On the other hand, there is an obvious mismatch between the complexity and ability of these systems to act in our world, and the level of knowledge required to interact critically with them. While the possibility to use them without expert knowledge has been a major factor in the diffusion of computing systems, it also has the important drawback that computing systems are nowadays used mostly by people who are unaware of the risks and consequences involved. Additionally, the increasing complexity and size of those systems, which is often rooted in a historically accumulated set of layers of abstraction and so-called bloated software systems (Wirth, 1995), has only deepened the issues of software design, development and maintenance as they came to be known in the 1960s. By consequence, it has become more difficult to prevent (potentially disastrous) errors. While this is principally a technical concern, it involves also political and commercial aspects underpinning the design, production and distribution of computing systems. From the point of view of the social and political implications, suffice here to mention issues of accountability in algorithm design and privacy of users.[2]

Given these circumstances, we are very much in need of a deeper reflection on the nature of computing systems. A methodological safe ground for such an investigation into the foundations of computing would require us to have a clear understanding of the field in itself, of the relations among its several sub-fields, and a solid grasp of how different approaches interact in the development of complex systems. There is, however, no such well-defined and clean foundation for the computing field in general: as it was argued in Tedre (2015), there is not even a clear and coherent identity. This is rooted, on the one hand, in the fact that computing has not yet reached its maturity as a discipline and, on the other, in that it is both a science and a technology, with often different and sometimes conflicting interests. While science aims at stable, durable and solid results, technology is driven by the need of quick innovation and even quicker market returns. As Kalanick, former CEO of Uber, has recently remarked in the context of discussions on self-driving cars:

> "*We are going commercial [...] This can't just be about science*".[3]

The DHST/DLMPST Commission for the History and Philosophy of Computing (www.hapoc.org) was established in 2013 with the awareness that such a fundamental reflection on the computing field can only be possible through the interaction and dialogue among different expertises. The approach of the Commission is

---

for some approach and overviews of the related literatures. For a high-level commentary on trust of digital technologies, see Taddeo (2017).

[2] The issue of algorithm accountability is gaining much traction, especially in view of current progress in AI. For a recent high-level analysis of the problem, see Diakopoulos, Friedler (2016). For contributions concerning the debate on the ethical relevance of algorithms in terms of accountability and their public impact, see Mittelstadt et al. (2016) and Binns (2017).

[3] Quoted in (Chafkin, 2017)

to create opportunities for collaborations and discussions within a pluri-disciplinary and pluri-methodological group of researchers, engaging with both the history, the philosophy and the formal and technical aspects of computing. We are strongly convinced that it is only by being embracive and tolerant with respect to different viewpoints, methods, and topics that it will be possible to develop a history and philosophy of computing which can account for both the scientific, social and technological aspects of the discipline.

Among others, one of the series of events organized by the Commission is the *Symposia on History and Philosophy of Programing* (HaPoP). The third in this series, HaPoP-3 was organized on June 25, 2016 at the *Conservatoire des Arts et des Métiers*, Paris by Liesbeth De Mol, Baptiste Mélès, Giuseppe Primiero and Raphaël Fournier-S'niehotta. Contrary to previous editions, this meeting focussed on one particular topic, namely on the nature, problems and impact of *operating systems*. The present volume collects contributions to HaPoP3.

Operating systems historically resulted from a broad set of general problems related to a large variety of aspects of computing (languages, memory, task complexity, to name a few) and so can be understood and contextualized as a (partial) answer to some of those problems. Moreover, both from the contemporary and historical perspective, it is hard to strictly isolate operating systems from others they are closely connected to, like networks and hardware systems. Accordingly, the editors have decided to shift the focus of the current volume to *programming systems*, to underline both the presence of historical aspects that precede and follow the birth of what qualifies as an operating system, the programming practices that underpin their design and development and the need to account for extensions of the concept of operating systems that we are witnessing today.

The general methodological approach of this book fits with the HaPoC philosophy. Accordingly, the current volume includes papers motivated by conceptual issues or questions alive in the contemporary debate but with roots in early episodes of the history of computing, and more historical contributions which bring to the fore technical problems still pressing today. The pluralistic approach of this book allows and even necessitates to overstep boundaries between communities and it is our hope that this effort will engage researchers to advance the very much needed multi-disciplinary foundations of computing.

2 Introducing Programming Systems

The term 'program' can have different meanings, and the historical context taken as a starting point for the origins of the activity of 'programming' largely affects the accepted definition. For the present purposes, we chose to start with ENIAC, since this is the historical context in which our contemporary use of 'program' originates.[4] ENIAC, one of the so-called first computers, was unveiled to the public in 1946. This

[4] See (Grier, 1996). Obviously, an ENIAC program is something quite different from a program expressed in a high-level language. See (Haigh, 2016b) for a different approach in which one starts

machine, in its initial configuration, had two fundamental properties that no other computing device had at that time:

1. it was an electronic machine, and so computation was done at a very high, humanly impractical speed;
2. it was programmable, i.e. it could be set-up to compute any function within the material limits of the machine.

It was this combination of high-speed and programmability that required a deeper reflection on both the design of the machine and the 'art' of programming it – two strongly connected aspects in ENIAC, where 'programming' meant physically rewiring the machine. First, there was a large gap between the time required to prepare and set-up a program and its execution time; second, unlike other contemporary machines like the Mark I, it was no longer possible to 'follow', and hence to fully control a computation; and finally, it made no sense to provide 'code' through the mechanically and slowly punched cards. The answers to these issues were twofold: first of all, ENIAC was permanently rewired as a stored-program machine and a new design, known generically as the EDVAC or von Neumann design, was described; secondly, different approaches for controlling "*the automatic evolution of a meaning*" (Goldstine, 1947) were developed.[5]

The ENIAC was a one-of-a-kind machine and by the late 1940s-early 1950s the standard design had stabilized on the EDVAC design (Neumann, 1945)[6] and the stored-program.[7] The latter is today considered as the stabilising technical and conceptual element from which 'programming systems' became possible: the stored-program concept expresses the basic principle of computer science that programs and data are interchangeable and granting, ultimately, the possibility to "*[simplify] the circuits at the expense of the code*" (Turing, 1946). As it will become clear from several contributions to this volume, contrary to what Goldstine and von Neumann believed, i.e. that:

> *"the problem of coding routines need not and should not be a dominant difficulty. [In] fact we have made a careful analysis of this question and we have concluded from it that the problem of coding can be dealt with in a very satisfactory way."* (Goldstine, 1946)

---

from a generic definition of 'program' (as a 'sequencing of operations'), and a 'modern program' is rooted in the ENIAC machine and the EDVAC design.

[5] More particularly, the approach taken by von Neumann was the identification of different steps in the preparation and set-up of a problem – a kind of division of labor – where the most prominent stage is that in which the 'dynamics' of a program is captured by means of a flowchart. The other is due to Curry, who focussed on the automation of the coding process and developed a logic for program compositionality. See (De Mol, 2015) for a partial comparison between the two approaches.

[6] It should be added that this is the standard narrative. Of course, there were many variants on the EDVAC design and also entirely different designs such as that for the Whirlwind which was not serial. See also (Backus, 1978) for a critical discussion of the von Neumann method.

[7] There are different understandings of the origins of the stored-program concept, its intellectual lineage and its historical implementation and understanding. See Haigh et al. (2014) and copeland-sommaruga for two different interpretations.

programming problems would not be resolved nor absolved by this basic principle, nor by initial symbolizations of the general flow of a program.

Once the design of computing machines had more or less stabilized, the construction of computing machinery moved away from the research labs at the university to industry, and so commercial interests started to play their role. However, for both scientific (e.g. SWAC) and business-oriented applications (e.g. in the context of LEO machines), and thus also for machines used in both contexts (IBM and Burroughs), there remained an important set of programming and physical problems to be resolved. Originally, computers were coded through machine instructions and so the semantic gap between "code" and hardware execution was quite small. However, this coding through stacks of punched cards or tape was a highly time-consuming and very error-prone process. Another associated issue was developing at the hardware level: the need of increasingly complex sets of instructions to execute meant the need for greater amounts of memory and required the ability to centralize the different instruction controls in one physical unit. In the 1950s, these two problems were tackled in a timely fashion and almost in parallel, as analysed in Part I of this volume.

First, one sees the development of techniques to optimize the coding process: while, at first, these were mainly developed and used in one particular practice and around a specific hardware system, there were clearly attempts at more systematic approaches. One well-known example is the programming book for the EDSAC (Wilkes, 1951) which is basically a 'library' of more or less standardized subroutines and, in its reprinted version, wanted to transcend the particularities of EDSAC to be a "*general introduction to programming for any computer of the stored-program type*". These approaches went hand-in-hand with improvements on the hardware design to have more efficient and simpler coding, with fewer errors or better error-handling.[8] The example of the LEO machines presented in Chapter I is illustrative of these initial, more systematic attempts. More particularly, it focuses on how approaches are developed to identify and resolve errors in a fashion that anticipates the definition of principles of correctness inspiring modern research in program verification. These efforts were strongly characterised by the business nature of the Lyons company who developed the machine: for this reason, more formal principles of valid program execution (like ensuring termination) are accompanied by heavily pragmatic choices (e.g. in the way program were designed and tests performed). Second, the improvement on hardware was essential not only in reducing the risks associated with component failures, but in particular in guaranteeing the possibility of accomodating a lot more memory. Complementary to these more systematic approaches that remain, basically, in the realm of the order, assembler and machine code, the first steps are being taken towards the development of higher-level

[8] See for instance the development of microprogramming which is basically an approach to hardware programming (Wilkes, 1953, p.230): "*This paper describes a method of designing the control circuits of a machine which is wholly logical and which enables alterations or additions to the order code to be made without ad hoc alterations to the circuits*" (De Mol, 2017) discusses several machines, some of which fit into the microprogramming strategy, that stick close to the hardware and develop optimum coding techniques such as latency and underwater programming.

languages and techniques of 'automatic coding' which were aimed at relieving the programmer from the tedious coding task and so to 'automate' the programmer.[9] For instance, Grace Hopper made the first steps in automating subroutining for the A0-language of UNIVAC I.[10] Chapter 8 is to be set against the background of such automating aspects of the programming process. It focuses on the period 1954-1964, the decade before the term 'operating system' more or less stabilized. This chapter deconstructs a classic narrative from the history of computing, that the operating system is essentially the IBM vision of automating the operator and hence it has to be historically located at the transition from batch processing to time-sharing systems. It is argued, instead, that this narrative in fact hides a more complex history concerning automating different aspects of programming. It is shown that it took several years before one could start to differentiate clearly between different kinds of systems, including the operating system. More particularly, a taxomomy of different types of systems from the late 50s and early 60s is offered. It is within this taxomomy that the steady development of 'operating systems' is accounted for, with the automation of the operator being just one of a set of parallel developments that brought about the distinguishability of operating systems from other ones.

The Complexity of Programming Systems

In the early 1960s, it is no longer mainly the hardware that shapes the problems related to the design and development of large-scale programming systems, but rather the programming systems themselves that determine the problems at stake. As summarized by the developers of the AOSP operating system for the Burroughs D825, (Anderson, 1962):

> "computers do not run programs, [...] programs control computers."

This underlies the diffuse realization that the problem does not lie with the machines, but rather with the 'programs', the way they are written, the way they constitute a complex system ready for commercial and scientific use by a broad range of different 'users' with different aims, who want well-documented, error-free and efficient systems. What is hard, is the software-side of computing.[11] From realizing that the programming problem was not simply going to be resolved by faster systems with larger memory, (the legend of) the software crisis was born. Dijkstra, in his Turing award lecture (Dijkstra, 1972), illustrated the situation as follows:[12]

> *"instead of finding ourselves in a state of eternal bliss with all programming problems solved, we found ourselves up to our neckys in the* software crisis *[m.i.]! [...] To put it quite*

[9] (Daylight, 2015), for instance, discusses the approach of the so-called 'space cadets'. See also (Nofre, 2014) for a discussion of the use of the notion of language in this context.

[10] See (Hopper, 1980) for a personal account of that development.

[11] See (Mahoney, 2008) for a historical take on this wordplay.

[12] See (Haigh, 2010) and (Ensmenger, 2010) for two different interpretations of the impact of the so-called software crisis.

> *bluntly, as long as there were no machines, progamming was no problem at all; when we had a few weak computers, programming became a mild problem, and now we have gigantic computers, programming has become an equally gigantic problem.[...] To put it in another way: as the power of available machines grew by a factor of more than a thousand, society's ambition to apply these machines grew in proportion, and it was the poor programmer who found his job in this exploded field of tension between ends and means."*

The 'software crisis' is thus closely tied with the development of progressively larger and more complex systems. Two approaches can be identified as answers to the problem:

1. the formal approach, closely associated with a logical understanding of the foundations of computing;
2. the modernist approach, associated with the development of 'grand designs' that would provide universal environments for the solution of multiple programming problems.

The contributions in Part II of this volume provide an understanding of the first of the two approaches above.

A very direct case of treating a programming system as a formal system is the development of programming semantics: these were introduced precisely to deal with issues related to mapping specification with implementation on possibly different machine architectures. Chapter II goes back to a basic historical case to discuss and compare four different styles of formal semantics that were developed in that context. In particular, the paper focuses on the origins and problems of four formal semantics developed for Algol 60: the Vienna operational description; Vienna functional description; Oxford denotational description and the VDM denotational description. It was considered a good language to demonstrate the potential of a logical approach, since it was supposed to scale well to realistic languages. For each of the four styles, the authors discuss not just the historical context, but engage with stylistic, syntactic as well as modelling features. This uniform approach for analyzing each of the styles not only results in a historical study of the reasons and modalities of their origins, but it also allows a more critical review of each semantical description of the language.

The formal approaches did not come about in a straighforward manner, but are based on a critical study of the foundations of programming and computing. Amongst others, they require an analysis of what constitutes a full-fledged computational object, including whether some 'objects' are going to have fewer rights than others, thus introducing a principle of non-uniformity. Chapter 7.7 engages with this problem, in particular with the possibility of treating functions as so-called first-class citizens. This basic problem in the foundations of programming, first pointed out by Strachey in the Algol context,[13] opens up a discussion on the technical and conceptual consequences of a particular formalization of computational citizenship; but it also connects developments in programming with the foundational debate in

[13] During a talk titled *Fundamental concepts of programming languages* given at the International Summer School in Computer Programming in Copenhagen, in August 1967, (Strachey, 1967)

mathematics from the late 19th and early 20th century. Within this setting, the nature of operating systems and their coming about is investigated at an even higher level of abstraction than language: the system becomes an environment in which functions can be treated as computational objects. This understanding of an operating system has the advantage of generating the conceptual space for a number of other associated (both theoretical and technical) topics fully developed in the modern understanding of systems: execution and access privileges, and the possibility of delegating them in different environments. It is especially interesting how establishing these traits of operating systems as environments of function definition and execution makes possible the convergence of both formal and technical discourse, in a move highlighting the foundation of both theoretical and physical computing.

The two contributions in Part III of this volume approach the mentioned 'modernist' take on the problem of system complexity. They both center their analysis on what is often considered one of the most successful of the grand design approaches to operating systems, namely Unix. In Chapter III, Unix is approached through its relations with other designs, both those that were supposed to improve it (like Plan 9), and those that had different philosophies (like Smalltalk). This analysis highlights a number of important features that have emerged as unifying traits in the process of system design: the focus on programmability as the main core-business of the system; the creation of a meta-system providing a unified semantic description for different types of objects (e.g. programs, files, devices); and its flexible 'everything is a file' design, allowing any program to be used with any file as input and any device as output. It is argued that both Unix and Smalltalk, while usually interpreted as 'grand designs', can also be aligned with a more postmodern understanding of programming in which there is not just one ultimate language but many, where each offers its own "viewpoint". The operating system as environment then becomes the backbone to support such postmodernist position.

In Chapter 9, Unix is set against the background of one of its predecessors, the Multics project: this detailed analysis of the features and processes in its early instantiation PDP-7 Unix, shows the switch from a 'bigger is better' approach to a 'simple is better'[14] one (Raymond, 2003):

> "Write programs that do one thing and do it well. Write programs to work together. Write programs to handle text streams, because that is a universal interface".

The process of creating the first Unix system started in 1969 and its several versions were developed until the 1980s, while the success of later instantiations like Linux and McOS are well-known to everyone.

3 Programming Systems in the Real World

The discussion whether a real software crisis has been overcome, or whether we are witnessing a new one, is still very alive today. Only, the stakes are now much

[14] Which is basically the 'worse is better' philosophy. See also Chapter III

higher: energy grids, banking systems, border controls, medical appliances, traffic control and automated vehicles, polling systems, any single important aspect of our everyday life is managed by and relies on a programming system. Part IV of this book engages with these more recent developments focusing on ethical, political and even aesthetical issues of large-scale systems.

In Chapter IV, the problem of defining ethical principles for operating systems within a safety-critical setting is analysed. Note how this project relies on the very same idea that motivates the formal approach to computing from the previous Section: first, the authors seek to connect processing in an ethical cognitive calculi to a successful, proof-based analysis and verification at the OS level; second, this formal analysis is implemented in a language to demonstrate feasibility in a self-driving system. The importance of logically grounded, verifiable and formally reliable systems like the present one is expanding and it is testified by its transfer from purely academic research to projects developed by major private players in the computing industry (Amazon and Facebook are particularly significative examples).

In Chapter 11 the relation between operating systems and globalization is examined from the point of view of sovereign nations, which are reconfiguring themselves as properly cyber-physical entities whose control extends to the software and data domains. The integration of state-sponsored and private software and hardware components is aimed at increasing control and at infringing user privacy: this aspect becomes nowadays essential in understanding the novel functional configuration of operating systems. The complexity of systems is thus again at stake in updating their definition, although this time with an additional level of influence, extending throughout the whole digital chain.

Finally, the definition of aesthetical criteria for complex systems relies inevitably on a compositional approach, almost matching the complexity analysis suggested by Fetzer (1988). In this respect, a first step is made in Chapter 7, where the problem of defining elegance of simple programs is tackled. This notion is analysed in terms of properties depending both on abstract and pragmatic criteria. These necessarily include the program's 'fitness for purpose', a criterium that (again) recalls the correctness principles mentioned at the very beginning of this volume.

The present volume is the first ever published combining historical, philosophical and technical approaches to tackle issues of programming systems. Whereas there are several studies focussing on single aspects and methodologies,[15] a combined approach, that allows to see different issues from multiple perspectives, was still missing from the literature. We consider this volume as a way to open up a very much needed foundational debate requiring the perspectives from historians, philosophers and practitioners: the former provide the historical backbone for current issues and so help distinguishing the 'real' issues from the more contingent ones; philosophers help discerning conceptual trajectories and the evolution of ideas, like those of correctness and computational citizenship; finally, the practitioners give the problem

[15] For historical works, see (Brennecke, 2002) and (Hashagen, 2002). For computer science works, see (Tanenbaum, 2008) and (Silberschatz, 2011); for a philosophical approach to software with some aspects related to systems, see Berry (2011).

context in which those ideas and issues originated, were technically tackled and evolved.

Gent, London *Liesbeth De Mol*
December 2017 *Giuseppe Primiero*

References

Anderson, J.P.; Hoffman, S.H.; Shiman, J and Williams, R.J. (1962). The D-825, a multiple-computer system for command & control. *1962 Fall Joint Computer Conference (AFIPS)*, pp. 86-96.

Backus, J. (1978). Can programming be liberated from the von Neumann style? A functional style and its algebra of programs. *Communications of the ACM*, vol. 21, nr. 8, pp. 613–641.

Berry, D.M. (2011). *The Philosophy of Software – Code and Meditation in the Digital Age*. Palgrave MacMillan.

Binns, R. (2017). Algorithmic Accountability and Public Reason. *Philosophy & Technology*. https://doi.org/10.1007/s13347-017-0263-5.

Brennecke, A. and Keil-Slawik, R. (1996). *History of Software Engineering*, August 26-30 1996, Dagstuhl seminar 9635, organized by W. Aspray, R. Keil-Slawik and D.L. Parnas

Chafkin, M. (2016). Uber's first self-driving fleet arrives in Pittsburgh this month. *Bloomberg Businesweek, August 18, 2016*. .

Copeland, B.J., Sommaruga, G. (2015). The Stored-Program Universal Computer: Did Zuse Anticipate Turing and von Neumann?, In G. Sommaruga, T. Strahm (eds.), *Turing's Revolution*, pp. 43–101, Springer International Publishing Switzerland.

Daylight, E.G. (2015) Towards a historical notion of 'Turing-the father of computer science'. *History and Philosophy of Logic*, vol. 36, nr.3, pp. 205–228.

De Mol, L.; Carlé, M.; Bullynck, M. (2015) Haskell before Haskell. An alternative lesson in practical logics of the ENIAC. *Journal of Logic and Computation*, vol. 25, nr.4, pp. 1011-1046, A version is available from: http://hal.univ-lille3.fr/hal-01396482/document

De Mol, L.; Bullynck, M.; Daylight. E. (2017). Less is more in the Fifties. Encounters between Logical minimalism and computer design during the 1950s. Available from: hal.univ-lille3.fr/hal-01345592v2/document

Diakopoulos, N., Friedler, S. (2016). How to Hold Algorithms Accountable. MIT Technology Review. https://www.technologyreview.com/s/602933/how-to-hold-algorithms-accountable/.

Dijkstra, E. W. (1972). The humble programmer, *Communications of the ACM*, **15**, 859–866.

Ensmenger, N. (2010). The computer boys take over, MIT Press.

Fetzer, J. H. (1988) Program Verification: The Very Idea, *Communications of the ACM*, 31(9):1048-1063.

Goldstine, H.H. and von Neumann, J. (1947) Planning and coding of problems for an electronic computing instrument Volume 2 of *Report on the mathematical and logical aspects of an electronic computing instrument*, part I,II and III, 1947-48. Report prepared for U. S. Army Ord. Dept. under Contract W-36-034-ORD-7481.

Goldstine, H.H. and von Neumann, J. (1946) On the principles of large-scale computing machines. in: Aspray, W. and Burks, A (eds.), *Papers of John von Neumann on computing and computer theory*, MIT Press, 1987, pp. 317–348.

Grier, D.A. (1996) The ENIAC, the verb 'to program' and the emergence of digital computers *IEEE Annals for the history of computing*, vol. 18, nr.1, pp. 51–55.

Haigh, T. (2010). Dijkstra's crisis: The end of Algol and the beginning of software engineering: 1968-72' in: *Workshop on the history of software, European styles*, Lorentz Center, University of Leiden.

Haigh, T., Priestley, M., Rope, C. (2014). Reconsidering the Stored-Program Concept. *IEEE Annals of the History of Computing*, vol. 36(1), pp. 4-17.

Haigh, T.; Priestley, M.; Rope, Cr. (2016). *Eniac in action. Making and remaking the modern computer*. MIT Press.

Haigh, T.; Priestley, M. (2016). Where Code Comes From: Architectures of Automatic Control from Babbage to Algol *Communications of the ACM*, vol. 59, nr. 1, pp. 39–44.

Hashagen, U; Keil-Slawik, R.; Norberg, A.L. (eds.) (2002) *History of Computing: Software Issues*. Springer.

Hopper, G. (1980). Keynote address Wexelblat, R.L. (ed.), *History of Programming Languages*, ACM Press, pp. 7–24.

MacKenzie, D. A. (2004). Mechanizing Proof – Computing, Risk, and Trust. MIT Press.

Mahoney, M. (2008). What makes the history of software hard? *IEEE Annals for the history of Computing*, vol. 30, nr.3, pp.8–18.

Mittelstadt, B.D., Allo, P., Taddeo, M., Wachter, S., Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, July-December 2016, pp. 1-21.

Naur, P. and Randell (eds.), B. (1968). *Software Engineering. Report on a conference sponsored by the NATO science Committee*, Garmisch, Germany, 7-11 October 1968 Available from http://homepages.cs.ncl.ac.uk/brian.randell/NATO/nato1968.PDF.

Neumann, P. G. (2017). Trustworthiness and Truthfulness are essential. *Communications of the ACM*, vol. 60, nr.6, pp. 1–3.

Nofre, D.; Priestley, M. and Alberts, G. (2014) When technology became language: the origins of the linguistic conception of computer programming, 1950-1960. *Technology and Culture*, vol. 55, nr. 1, pp. 40–75.

Primiero, G., Boender, J. (2017). Managing Software Uninstall with Negative Trust. In Jan-Philipp Steghöfer and Babak Esfandiari (eds.), *Trust Management XI - 11th IFIP WG 11.11 International Conference, IFIPTM 2017*, IFIP Advances in Information and Communication Technology, vol. 505, pp.79–93, Springer.

Primiero, G., Raimondi, F., Chen, T., Nagarajan, R. (2017). A Proof-Theoretic Trust and Reputation Model for VANET. In 2017 IEEE European Symposium on Security and Privacy Workshops, EuroS&P Workshops 2017, pp.146–152.

Raymond, E.S. (2003). *The art of Unix programming*. Addison-Wesley Professional.

Silberschatz, A.; Galvin, P.B.; Gagne, G. (2011). *Operating system concepts* Wiley and Sons.

Strachey, C. (1967). Fundamental Concepts in Programming Languages. *Higher-Order and Symbolic Computation*, 2000, vol. 13, pp. 11-49.

Taddeo, M. (2017) Trusting Digital Technologies Correctly. *Minds & Machines*. https://doi.org/10.1007/s11023-017-9450-5.

Tanenbaum, A.S. (2008) Modern Operating Systems. Pearson International Edition, 3rd. edition.

Tedre, M. (2015). *The science of Computing. Shaping a discipline*. Boca Rato: CRC Press.

Turing, A. M. (1946) Lecture to the London Mathematical Society on 20 febru- ary 1947. 1947. in: Brian E. Carpenter and Robert W. Doran (eds.), *A.M. Turing's ACE Report of 1946 and Other papers*, MIT Press, 1986, 106-124.

von Neumann, J. (1945) *First draft of a report on the EDVAC*, University of Pennsylvania, June 30, 1945. Available from: http://www.virtualtravelog.net/entries/2003-08-TheFirstDraft.pdf.

Wilkes, M.V.; Wheeler, D.J. and Gill, S. (1951) *The preparation of programs for an electronic computer* Addison-Wesley Second edition, 1967.

Wilkes, M. and Stringer, B. (1953) Micro-programmingand the design of the control circuits in an electronic digital computer *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 49, nr. 2, pp. 230–238.

Wirth, N. (1995) A plea for lean software *Computer*, vol. 28, nr. 2, pp. 64–68.

Acknowledgements

We would like to thank Baptiste Mélès and Raphaël Fournier-S'niehotta for their help with setting up and chairing HaPoP-3. We are also very grateful to the participants to the symposium as well as the PC members for their help in selecting the accepted talks. Finally, this book volume would not have been possible without the careful and critical reading of the different reviewers who helped to improve the contributions.

Part IV
Evaluating Systems

Ethical Operating Systems

Naveen Sundar Govindarajulu, Selmer Bringsjord, Atriya Sen, Jean-Claude Paquin, Kevin O'Neill

Abstract A well-ingrained and recommended engineering practice in safety-critical software systems is to separate safety concerns from other aspects of the system. Along these lines, there have been calls for operating systems (or computing substrates, termed *ethical operating systems*) that implement ethical controls in an ethical layer separate from, and not amenable to tampering by, developers and modules in higher-level intelligence or cognition layers. There have been no implementations that demonstrate such a marshalling of ethical principles into an ethical layer. To address this, we present three different tracks for implementing such systems, and offer a prototype implementation of the third track. We end by addressing objections to our approach.

Naveen Sundar Govindarajulu
Dept. of Cognitive Science, Dept. of Computer Science, Rensselaer AI & Reasoning (RAIR) Lab, RPI, Troy NY 12180 USA, e-mail: naveensundarg@gmail.com

Selmer Bringsjord
Dept. of Computer Science, Dept. of Cognitive Science, RAIR Lab, Lally School of Management, RPI, Troy NY 12180 USA, e-mail: selmerbringsjord@gmail.com

Atriya Sen
RAIR Lab, Dept. of Computer Science, RPI, Troy NY 12180 USA, e-mail: atriya@atriyasen.com

Jean-Claude Paquin
RAIR Lab, RPI, Troy NY 12180 USA, e-mail: paquij@rpi.edu

Kevin O'Neil
RAIR Lab, RPI, Troy NY 12180 USA, e-mail: oneilk4@rpi.edu

1 Introduction

Suppose that r is an intelligent, autonomous, robot whose range of human-impacting actions in the environment is wide and substantive. Govindarajulu and Bringsjord (2015) have explained and defended, at length, the following two-part position:

P$_1$ r will need to be ethically controlled; and

P$_2$ such control cannot be achieved by merely installing high-level modules that monitor the ethical status of r's actions, but rather only by infusing the OS of r with computational logics of the right sort (see Figure 1).

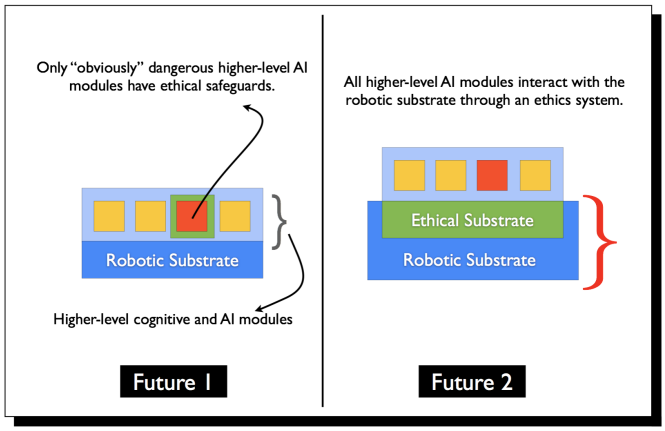

Fig. 1 Two Futures — With and Without an Ethical Substrate *Higher-level modules are vulnerable to tampering. The Ethical Substrate protects the Robotic Substrate from rogue modules*. Figure from Govindarajulu and Bringsjord (2015)

To ease exposition, we assume that **P**$_1$ is granted. Then the main rationale for **P**$_2$, encapsulated, is this: Unless ethical control is engineered at the operating-system level, malevolent or blundering software engineers working above the OS level may well disable such control. There is a simple software-engineering-motivated rationale for needing ethical operating systems, as shown in Figure 2. By offloading the development and refinement of ethical theories,[1] AI developers can focus on building intelligent systems and need not be concerned with the esoteric ins and outs that

[1] We here use the word 'theory' as it is used in formal logic and mathematics; there, a *theory* is any arbitrary set of formulae Γ (which may e.g. be the closure under deduction of some set of core axioms). Hence, for us, an *ethical* theory is a set of formulae that governs ethical behavior. Coverage of such theories ranges from the simple, such as a list of prohibitions, to the more complex, e.g. the *doctrine of double effect* (discussed herein later), and beyond. Our conception of an ethical theory is in the end simply a rigorization of the concept of an ethical theory as employed by analytic ethicists, an examplar being Feldman (1978); a synoptic explanation of this is given in Note 10. Our sense of 'ethical theory,' then, is in the end a formal version of what systematic ethicists refer to when they discuss such ethical theories as utilitariaism, ethical egoism, contractualism, etc.

are the bread and butter of professional philosophers and other experts. This philosophical work can be assigned to those trained for such work. This approach can be seen as an application of the principle of *separation of concerns* in Dijkstra (1982).[2]

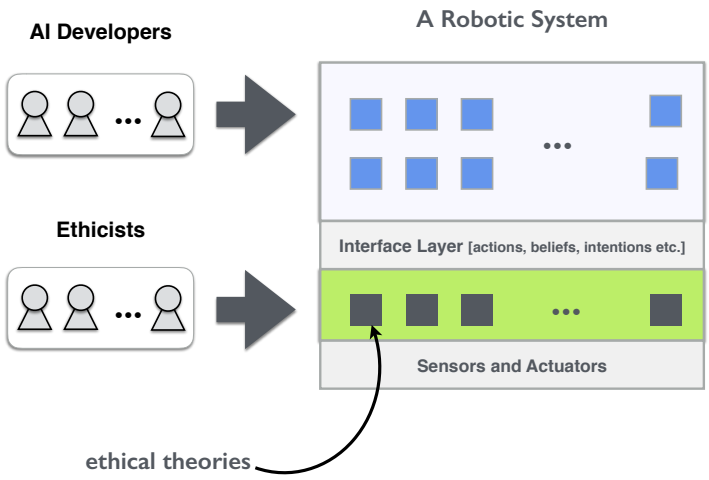

Fig. 2 The Goal: Software-Engineering Perspective on an Ethical Operating System

In other, directly related prior work, Bringsjord and Sen (2016) have made the sustained case that, where *r* is specifically a self-driving car, OS-rooted ethical control on the strength of the right sort of computational logics is necessary (despite what sanguine car manufacturers may currently believe). Unfortunately, while we claim to have in hand the required computational logics for ensuring that when possible *r*, relative to some selected ethical theory, meets all its moral and legal obligations, never does what is morally or legally forbidden, invariably steers clear of the invidious, and, when appropriate, performs what is supererogatory,[3] to this point

[2] It is quite easy to see how Dijsktra's principle still applies when we want to engineer ethical machines, for we read:

> We know that a program must be correct and we can study it from that viewpoint only; we also know that it should be efficient and we can study its efficiency on another day, so to speak. In another mood we may ask ourselves whether, and if so: why, the program is desirable. But nothing is gained — on the contrary! — by tackling these various aspects simultaneously. It is what I sometimes have called 'the separation of concerns,' which, even if not perfectly possible, is yet the only available technique for effective ordering of one's thoughts, that I know of. (Dijsktra 1982, p. 60)

[3] One calculus that enables much of this is the **deontic cognitive event calculus** (with provision for modeling access/informational self-awareness), or for short $\mathcal{DCEC}^*$, which has now been used

we have not worked directly at the operating-system level in any detail, and *a fortiori* have no demonstration that OS-rooted ethical control of *r* can be specified and implemented. In the present contribution, we lay out a **formal meta-operating system** and describe an embryonic implementation of it that carries a non-trivial ethical component. We also end by entertaining and rebutting some penetrating objections to our "meta" approach.

2 Prior Work in Ethical Control

We plan to be able to concretely demonstrate not only that our ethical-control calculi can ensure that the robots meet their obligations to, for instance, protect life (an example of which is shown in Figure 3, where Bert from Sesame Street is saved in the RAIR Lab from being run over by an onrushing car when the saving car deflects the onrushing one), but that such morally correct behavior can be OS-rooted. In the present section, however, we say a few words about prior work in ethical control of robots, *simpliciter*.

There are more than a few projects for ethical control of robots based on logic-based/logicist formalisms. The *Deontic Cognitive Event Calculus**, $\mathcal{DCEC}^*$, a quantified multi-operator modal logic, has been used to model not only obligatory actions like saving Bert by deflection (again, Figure 3), but also for example *akrasia* (willful violation of one's own self-affirmed moral principles, Bringsjord et al. (2014)), and the doctrine of double effect (Govindarajulu and Bringsjord, 2017).

In addition, Pereira and Saptawijaya (2016a) use a propositional logic programming approach to model not only the doctrine of double effect, but many other phenomena relevant to — as they put it — "programming machine ethics" (Pereira and Saptawijaya, 2016b). In addition, since any mechanization of explicit laws or principles that preserves their declarative content in symbolization that is reasoned over classically is fundamentally logic-based, much of the early, seminal work of Arkin (2009) is by definition in the logicist paradigm. Additional early machine-ethics work that is explicitly logicist includes Arkoudas et al. (2005); Bringsjord et al. (2006). And, to mention a final example of prior research, in some very important work based in answer-set programming, Ganascia (2015) has tackled the problem of using non-monotonic logic to model and resolve conflicts in ethical reasoning.[4]

None of the work referred to in the previous paragraph, please note, is connected to OS-level processing in any way; the same holds for research in the same vein that we don't explicitly cite. If $\mathbf{P}_2$ holds (refer to the beginning of Section 1), then this is undesirable. This is simply an observation, one devoid of any criticism of the

in its implemented form to guide and control the actions of a number of real-life versions of what *r* denotes in the present paper; e.g. see Bringsjord et al. (2014). The earliest work of this kind started over a decade ago (Bringsjord et al., 2006; Arkoudas et al., 2005), and has been steadily improving — but hitherto has not been connected to operating systems. An overview of $\mathcal{DCEC}^*$ can be found at this url: http://www.cs.rpi.edu/∼govinn/dcec.pdf.

[4] See also the earlier Ganascia (2007).

intrinsic quality of the work itself; note that the observation is accurately made of prior work in our own case. We turn now to two straightforward "tracks" that can be pursued.

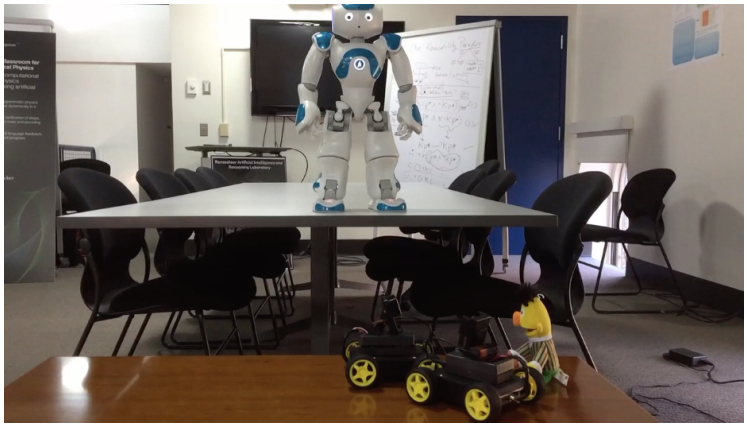

Fig. 3 A Demonstration of *Obligation-only* Ethical Control *The self-driving robot to the left of Bert would have run him over — but the other self-driving robot met its obligation by deflecting the onrushing car, thereby keeping Bert and his acting career alive and well. The robot overhead on the table is ethically controlled as well, but realized that it didn't have an obligation to dive down to save Bert.*

3 Two Possible Tracks

There are two possible tracks that naturally come to mind when one is looking to achieve an ethical operating system. Track 1 is aimed at realistic-scale, purely formal vindication of our approach to ethical operating systems. Here, in our own case, we would seek to connect processing in our ethical cognitive calculi to successful, real-world proof-based analysis and verification at the OS level.[5] In Track 1, our ethical-control logics would be interleaved with seL4 to form what Govindarajulu and Bringsjord (2015) dub the *ethical substrate*, and the goal would be to establish this at the conceptual/formal level first, before moving on to implementation. By "interleaving" an OS with an ethical calculus, we mean: (1) the combination of any

[5] At the moment, among formally verified operating-system kernels, the clear frontrunner is apparently seL4 https://sel4.systems. It runs on both x86 and ARM platforms, and can even run the Linux user-space, currently only within a virtual machine. It's also open-source, including the proofs. These proofs can be combined with our own for ethical control. For a remarkable success story in formal verification at the OS-level, and one more in line with the formal logics and proof theories our lab is inclined to use, see Arkoudas et al. (2004).

formal calculus and theory used in the verification of the system with the ethical calculus; and (2) use of the ethical calculus in the OS during its operation.

Track 2 is much more concrete; in it, we are working in what can be called "microcosmic" fashion, leveraging theorem proving and a formalization of a subset of Common Lisp. Here we are building a miniature operating system for mobile robots that run our ethical-control calculi, to regulate and control the behaviour of the system. We are seeking to include these calculi in this system so as to demonstrate feasibility in the self-driving-car domain.[6] We are doing this for miniature self-driving cars, and a key part of our work is the use of ACL2.[7] See Figure 4.

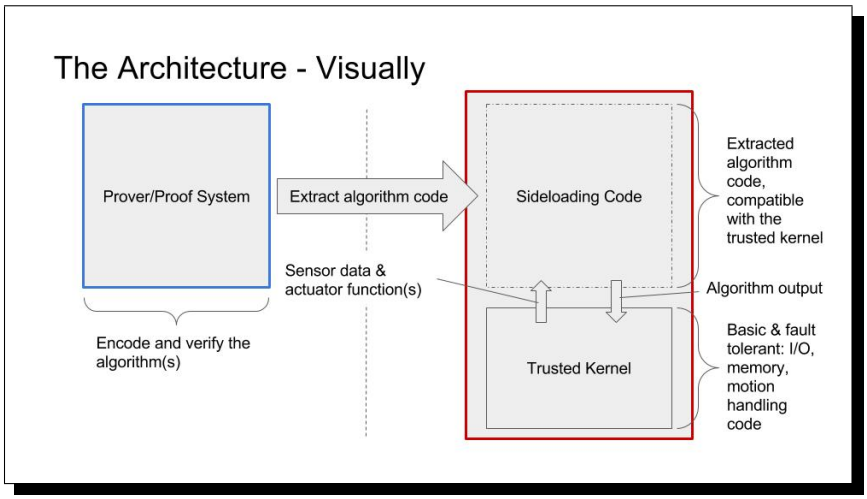

Fig. 4 An Architecture for a Mobile Robot OS *The proof system is ACL2.*

4 Track 3: A Blend of Tracks 1 and 2

We now move to the technical focus of the present paper, in the context of our foregoing synopses of Tracks 1 and 2: viz., a hybrid track that marks a "blending" of these two tracks. This blended approach we refer to as 'Track 3.' The rationale for adding Track 3, and pursuing it, is fairly straightforward. This rationale begins by conceding a brute fact: Engineering an operating system from the ground up, a là Track 2, even when the range of coverage for the computation in question is severely

[6] At least at the conceptual level, there is some historical precedent for at least the first steps of what we are seeking: Flatt et al. (1999) showed that "MrEd," while not a "bare-metal" OS, is a Lisp-flavored virtual machine that counts as an OS.

[7] 'ACL2' abbreviates 'A Computational Logic for Applicative Common Lisp.' The home page is: http://www.cs.utexas.edu/∼moore/acl2.

restricted, is a gargantuan task. At the same time, however, the formal rigor of Track 1 must be conceded to be attractive, and the prospect of connecting work on ethical operating systems to the longstanding, excellent, and rich body of methodologies and work on program verification is a very savory one.[8] Track 3, if you will, enters this situation and "comes to the rescue." We are still pursuing Tracks 1 and 2, but Track 3 is what we emphasize in what follows, since it allows us to quickly make advances worth (at least by our lights) sharing with readers. The fact is that up until now, all published work by us in the domain of ethical operating systems has been abstract, and at the same time, all of our engineering work has been exclusively in machine ethics, divorced from connections to operating systems.

In Track 3, instead of engineering an operating system from the ground up or building a simple, formally-verified kernel, we look at building an ethical **meta-operating system**. A meta-operating system is a system that runs on top of an existing operating system, yet provides all the routine functions of an operating system (such as managing hardware) to software that in turn runs on top of *it*. We begin by extending our prior work in this arena by specifying a formal model for a meta-operating system. In order to do that, we first need to make more precise a few common and useful concepts. The first of these is the notion of *software components*:

Software Components (Abstract)

We begin by assuming as primitives a set of all possible *software components* $\mathcal{S}$. Any robotic or computational system S, at any time $t \in \mathbb{N}$, has an associated finite directed graph $S_G(t)$ with nodes $S_N(t)$ and edges $S_E(t)$, with nodes $n \in \mathcal{S}$. An edge (u, v) indicates that component u is dependent on v. (See Figure 5.)

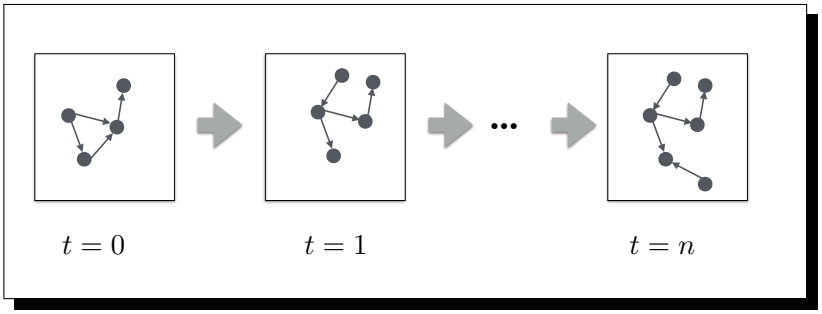

Fig. 5 A Software System in the Abstract

By 'software component,' we mean a running software process with internal states and not simply the definition or program that spawned the process. Armed with the above defintion, we obtain the following straightforward view of what an operating system is:

[8] For summary and references, see Bringsjord (2015b), which includes a defense of a particular way to seek verification.

Operating System (Abstract)

Given a system S, an *operating system* is simply the only unique component o such that for all times t, there is a path from any component $v \neq o$ to o. A path from to u to v is a sequence of one or more edges $\left[(u,p_1),(p_1,p_2),(p_2,p_3),\ldots,(p_{n\text{-}1},p_n),(p_n,v)\right]$.[a]

[a] In distributed systems, there can be multiple such components.

The definition immediately below states that a meta-operating system m is a software component dependent on another component, the underlying operating system o; and the rest of the components are transitively dependent on the meta-operating system m. In other words, the meta-operating system is simply another software component that sits between an operating system and all other components in a systen. ROS (the Robot Operating System) and Player/Gazebo (Vaughan et al., 2003) in the robotics domain are two such prominent meta-operating systems. Intuitively, a meta-operating system can be thought of as an interface to another underlying operating system.[9]

Meta-Operating System (Abstract)

Given a system S, for all times t, a *meta-operating system* is a software component m such that there is a component o (the underlying operating system that m uses) such that:

$$(m,o) \in S_E(t) \text{ and } \forall o' \cdot \exists t' \cdot (m,o') \in S_E(t') \rightarrow o = o'$$
$$\text{but for all } v \neq o \text{ and } v \neq m, \text{ there is a path from } v \text{ to } m$$

Though meta-operating systems such as ROS and Player/Gazebo differ quite a bit, the above semi-formal definition roughly captures the intended notion. As we mentioned above, though there have been calls for ethical operating systems and arguments for why such systems are needed, there has been very little work in either formal or real systems. In the rest of the paper, we present an ethical meta-operating system accompanied by an implementation. While the system is quite simple, it is concrete and available for researchers to experiment with and extend. We have the following informal definition for what constitutes an ethical operating system:

Ethical Operating System *Informal* Requirement

An ethical operating system $\mathcal{E}_E$ is an operating system that adheres to an ethical theory E even when software components are added, removed, or when configurations between components change.

We are well aware of the fact that 'adheres to an ethical theory E' is a broad phrase. However, since the focus in the present essay is specifically on presenting the (Track 3) conception of an ethical operating system, we leave aside here the fleshing out of this broad locution. Also, more precisely, our approach works with not just adherence to a given ethical theory, but adherence to ethical *codes* derived from a given theory.

[9] The definition that immediately follows does not distinguish between virtual operating systems and meta-operating systems and does not account for nested meta-operating systems.

In fact, our process overall consists in the four steps shown in Figure 6. Obviously the present chapter centers around Step 3: bringing machine ethics to OS-level processing.[10]

5 Ethical Calculi

There are a number of families of ethical theories. For example: deontological theories, utilitarianism, divine-command theories,[11] contractualism, virtue ethics, "ethical egoism," etc. (these are pictured schematically in Figure 6). We do not want to advance a framework that requires one to commit to any particular one of these theories or even to families of theories. Our framework is general enough that it can be applied to *any* ethical theory, or collection or family thereof. That said, there are a few high-level requirements that should be discussed and affirmed.

Assume that we have a family of ethical theories **E** of interest. We assume that any ethical theory $E \in \mathbf{E}$ obligates or permits (i.e. sanctions) a set of situations or actions Π and forbids a set of other situations or actions Υ. Any formal system in play must have enough power to capture these notions.

Abstractly, assume that we have a formal system $\mathcal{F} = \langle \mathcal{L}, \mathcal{I} \rangle$ composed of a language $\mathcal{L}$ and a system of inference schemata (or a proof theory/argument theory) $\mathcal{I}$. This system could be as sophisticated as $\mathcal{DCEC}^*$, a quantified multi-modal logic used, for example, in Bringsjord et al. (2014), or it could be as simple as *standard deontic logic*, a propositional modal logic, used in Govindarajulu and Bringsjord (2015). The only requirement is that the system be sophisticated enough to model any situation and condition the selected family **E** of ethical theories might have to

[10] While the focus of the present paper is on Step 3, we provide a brief explanation of the mysterious-to-most-readers phrase "run through EH" that appears in the graphic of Figure 6: An ethical theory T in the four-step process is formalized as a conjunction of robust biconditionals $\beta(x_1, \ldots, x_k)$ that specify when actions, in general, are obligatory (and forbidden and morally neutral); here, x_i are the variables appearing in the biconditional, and serve the purpose of allowing for the fixing of particular times, places, and so on. The general form of each definiendum of each biconditional refers to some action being $\mathcal{M}$ for some agent in some particular context; the definiens then supplies the conditions that must hold for the action to be $\mathcal{M}$. This is a rigorization of the approach to pinning down an ethical theory taken e.g. in Feldman (1978). The variable $\mathcal{M}$ is a placeholder for the basic categories captured by modal operators in our calculi. For instance, $\mathcal{M}$ can be *obligatory*, or *forbidden*, or *civil*, etc. Now, the ethical hierarchy $\mathscr{EH}$ introduced in Bringsjord (2015a) explains that this trio needs to be expanded to *nine* different deontic operators for $\mathcal{M}$ (six in addition to the standard three of *forbidden*, *morally neutral*, and *obligatory*). (For example, some actions are right to do, but not obligatory. A classic example is the category of *civil* actions. There are also *heroic* actions. The expansion of deontic operators to cover these additional categories was first expressed systematically in Chisholm (1982).) To "run a given ethical theory through $\mathscr{EH}$" is to expand the activity of Feldman (1978), for a given ethical theory, to biconditionals $\beta(x_1, \ldots, x_k)$ *for each of the nine operators*. (Feldman only considers one.) A particular code C_T based on an ethical theory T, if configured in keeping with $\mathscr{EH}$, would include use of any of the operators in the nine in order to e.g. permit or proscribe a particular kind of action in a particular domain for a given agent under T.

[11] Yes, even this family can be used for machine/robot ethics; see e.g. Bringsjord and Taylor (2012).

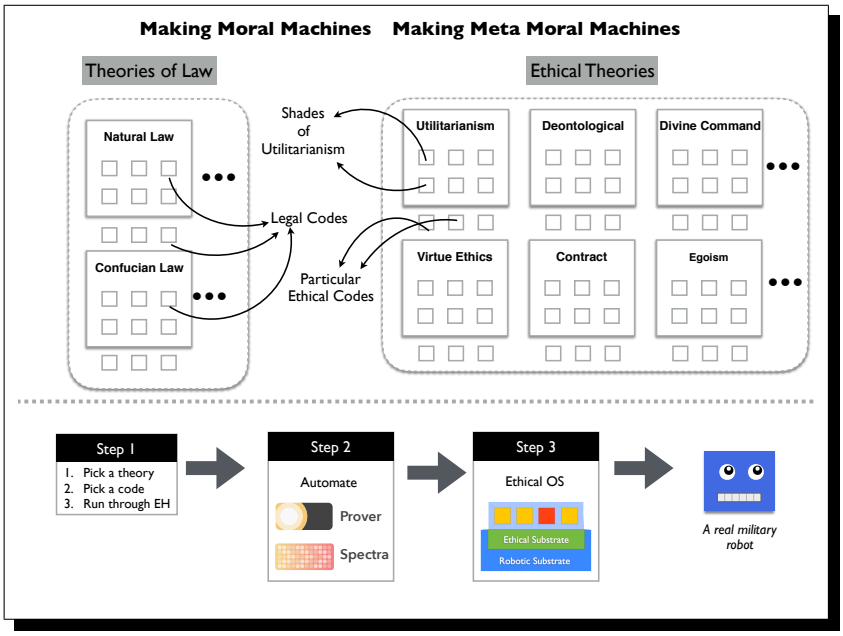

Fig. 6 The Four Steps in Making Ethically Correct Machines. Step 3, in broad strokes the connecting of mechanized ethics to OS-level processing, is the focus of the present chapter. For an overview of the four-step process, including some explantion of the 'Run through' sub-step in Step 1, see footnote 10.

handle. The requirement for our formal system $\mathcal{F}$ is that it has to be expressive enough to capture any theory $E \in \mathbf{E}$ via a set of formulae Γ_E in $\mathcal{L}$. We require that, for any sanctioned situation in $\pi \in \Pi$, there is a formula ϕ_π; and, for any forbidden situation $\upsilon \in \Upsilon$, there is a formula ϕ_υ representing it. With these requirements met, the following obvious conditions arise:

$$\Gamma_E \vdash_{\mathcal{I}} \phi_\pi$$
$$\Gamma_E \vdash_{\mathcal{I}} \neg\phi_\upsilon$$

We also need two more technical conditions to hold:

(1) For any given theory E, if E is *sound*, we require that Γ_E be *consistent*; that is, there is no ϕ such that $\Gamma_E \vdash_{\mathcal{I}} \phi$ and $\Gamma_E \vdash_{\mathcal{I}} \neg\phi$.
(2) Γ_E is *negation-complete*; that is, for any ϕ: $\Gamma_E \vdash_{\mathcal{I}} \phi$ or $\Gamma_E \vdash_{\mathcal{I}} \neg\phi$.

6 A Formal Meta-Operating System

We use the **actor calculus** to provide a model of a meta-operating system. The actor calculus is a Turing-complete model of computation used for modeling and building concurrent computing systems.[12] The actor calculus is well-suited for systems in which components have to be added or removed, and in which connections between components can change through time.

At the core of the actor calculus is — unsurprisingly — an *actor*, simply an independent unit of computing. In any computing system, there can be zero or more actors, each operating independently and concurrently. Actors communicate by exchanging messages. Each actor can be thought of as a "black box."

We now give a quick, semi-formal conceptualization of the actor calculus. Assume that we have a formal system $\mathcal{F} = \langle \mathcal{L}, \mathcal{I} \rangle$ as discussed above. Let N be a set of identifiers or names. We employ the simply typed λ-calculus and augment it with the following primitive expressions: $\{send, new, ready\}$, giving us the λ^a-calculus with which we construct actors. (The new primitives will be explained shortly.) Also assume that the set of expressions in the λ^a-calculus includes $\mathcal{L}$.[13] Let $\mathcal{B}$ be the set of all expressions of λ^a-calculus.

Actor Calculus Components (Modified)

Actor An actor (as stated earlier) is an independent unit of computation. Each actor has a unique name $n \in N$. An actor is associated with a λ abstraction (i.e., function definition) in λ^a-calculus.

Message A message is an element of $\mathcal{L}$.

The new primitives are explained immediately below.

Actor Calculus Components (Modified)

1. $send: N \times \mathcal{L} \times N \to \{\}$. $send(x, m, y)$ is used for sending a message m to an actor x from the actor y.
2. $new: \mathcal{B} \to N$. This primitive is used for creating a new actor with behavior specified by the input λ^a-calculus expression. The primitive generates a brand-new identifier for the actor.
3. $ready: \mathcal{B} \to \{\}$. This changes the invoking/calling actor's behavior to one specified in the input. This is useful for modeling components that change their internal state.

Note that while the above model is functional in nature, there are models of the actor calculus that use other programming paradigms. For instance, SALSA is a standalone actor-calculus-based programming language that runs on the Java Virtual Machine (JVM) and is object-oriented in nature (Varela and Agha, 2001). Akka is another JVM-based object-oriented actor system available as a library for languages

[12] In concurrent computing, there can be two or more different computational processes happening at the same time.

[13] The inclusion of an arbitrary formal language $\mathcal{L}$ is where we differ from the strict λ^a-calculus as presented in, for instance, (Varela, 2013, Chapter 4). This is merely for convenience and doesn't sacrifice generality, as we can readily encode $\mathcal{L}$ using primitives in just the λ-calculus and nothing more.

on the JVM (Boner, 2010). Our implementation of an embryonic ethical operating system uses an object-oriented framework based on Akka. For a purely functional system, see the cl-actors system for Common Lisp (Govindarajulu, 2010).

Defining Dependency

In the actor model, software components are actors. An actor u is dependent on an actor v *iff* the definition for u has the identifier for v.

We get our formal model of a meta-operating system by taking the most general description of an intelligent agent as can be found in Russell and Norvig (2009) and Hutter (2005) and casting that in an actor-based formalism. (These works incontestably provide supremely general accounts of what an intelligent agent is.) See Figure 7. We make the architecture shown in Figure 7 more specific by requiring that sensors, actuators, and the agent be composed of one or more actors. See Figure 8. Given the actor formalism, decomposing an agent architecture into actors is quite simple. We require that there be four classes of actors, or correspondingly four classes of names, as given below:

Actor-Calculus Agent System

Sen Names of actors that are used as sensors. These actors get information from the external environment.

Int Names of actors that are used as internal components. These actors perform the reasoning and any other cognitive tasks (learning, planning, and so on).

Act Names of actors that are used as actuators. These actors change things in the environment.

Env Names of actors in the environment. There could be just one actor modeling the entire external environment, or there could be a set of actors modeling different parts of the environment.

In the actors world, the operating system is then simply the collection of actors $\mathbf{Sens} \cup \mathbf{Act}$, as all other actors would need to transitively rely on these actors for interactions with environment.[14] Given this, a meta-operating system is then simply $\mathbf{Sens} \cup \mathbf{Act}$, or a fully encapsulating layer around $\mathbf{Sens} \cup \mathbf{Act}$.

7 A Formal *Ethical* Meta-Operating System

Since messages between actors are all from $\mathcal{L}$, specifying an ethical operating system becomes straightforward. At a minimum, we simply need all messages from any actor in **Act** to any actor in **Env** to be sanctioned by the ethical theory we are using. (Please recall our remarks in §4 in which we conceded that directly using an ethical theory is a gross simplification, but expedient given the current chapter's focus; and specifically recall the four steps alluded to in Fig. 6.) In the following condition, the ethical layer acts a filter or a gate. Under the pass condition, it lets the

[14] We ignore stray actors that neither observe nor act upon the environment.

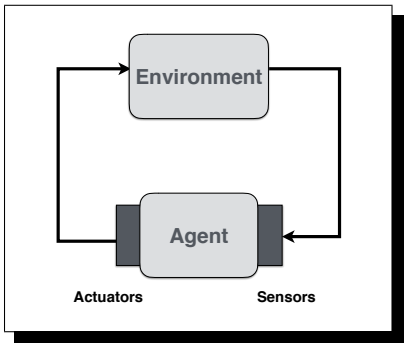

Fig. 7 Architecture for an Intelligent Agent

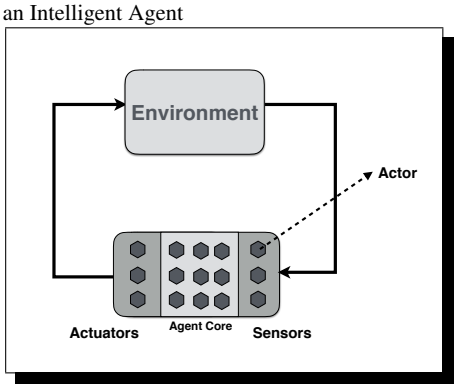

Fig. 8 An Actor-based Architecture

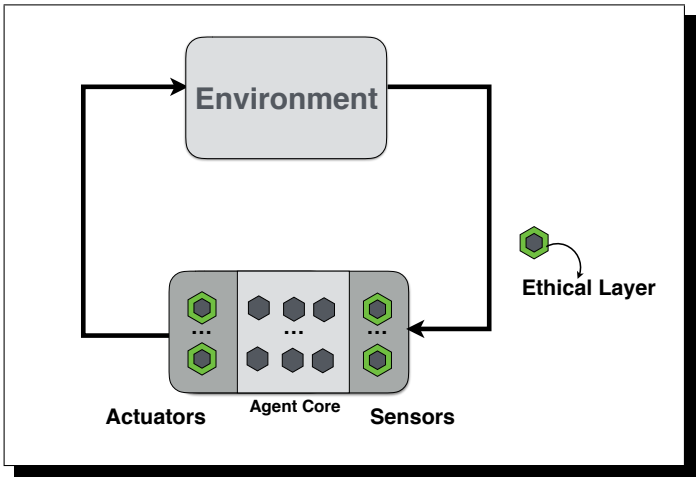

Fig. 9 Ethical Layer around the Meta-Operating System

message through; and under the fail condition, it simply discards the message (see Figure 9).

$$\forall q \cdot \forall r \in \mathbf{Env} \cdot \forall m \in \mathcal{L} \cdot \left[send(r,m,q) \Rightarrow \left(\textbf{if } E \vdash_{\mathcal{I}} m \textbf{ then } \mathsf{pass} \textbf{ else } \mathsf{fail} \right) \right]$$

The above condition works well for ethical theories that are only concerned with the actions of an agent; but the condition will fail when we rely on ethical theories that pivot on the internal, intensional states of agents. For example, the **Doctrine of Double Effect** ($\mathcal{DDE}$) requires considering one's intentions when weighing actions that have both good and bad effects.[15] Modeling the doctrine requires modeling an agent's knowledge and intentions. This requires us to consider internal messages too.[16] The condition becomes simpler to write but more expensive to check during the system's operation:

Ethical Layer Condition 1

$$\forall q \cdot \forall r \cdot \forall m \in \mathcal{L} \cdot \left[send(r,m,q) \Rightarrow \left(\textbf{if } E \vdash_{\mathcal{I}} m \textbf{ then } \mathsf{pass} \textbf{ else } \mathsf{fail} \right) \right]$$

The above two conditions look at only messages that have been sent and check whether they conform to the theory E or not. The conditions don't account for circumstances in which E dictates that a certain message has to be sent, but in fact no message is sent. Let the statement *"m should be sent to r at time t"* be denoted by the formula $\sigma(m,r,t)$. Then the layer, denoted by the actor l, can send such a message on its own if it confirms that no such message exists at t:

Ethical Layer Condition 2

$$\forall r \cdot \forall m \cdot \forall t \cdot \begin{pmatrix} E \vdash_{\mathcal{I}} \sigma(m,r,t) \wedge \neg \exists q \cdot send(r,m,q) \\ \Rightarrow \\ send(r,m,l) \end{pmatrix}$$

The above formulation gives rise to an immeditate concern. While the formulation constrains individual messages, the messages themselves can be at any level of abstraction and need not be just individual atomic actions that an agent might commit. For example, consider a prohibited action a composed of two or more actions $\langle a_1, a_2, \ldots, a_n \rangle$. The layer can correctly work in this case if a is sent as a message. If a is not sent as a message, the layer can keep track of $\langle a_1, a_2, \ldots, a_{n\text{-}1} \rangle$ and prohibit a_n, and thus prevent a from being realized.

[15] A rapid, informal, but nonetheless nice overview of the doctrine is provided in McIntyre (2014).

[16] A quick note on the expressivity of the formal system needed to model $\mathcal{DDE}$: It is well known that modeling knowledge in first-order logic can lead to fidelity problems by permitting inconsistencies. We show this explicitly in Bringsjord and Govindarajulu (2012). This implies that $\mathcal{DDE}$ requires going beyond first-order logic to first-order *modal* logic (an intensional logic) with operators covering minimally the epistemic and deontic realms. An intensional model of $\mathcal{DDE}$ can be be found in our Govindarajulu and Bringsjord (2017).

8 Implementation & Walkthrough

We now explain an embryonic implementation of a meta-operating system that has the facility to bake in ethical theories that operate independently of other modules in the system. This is achieved by implementing an ethical layer that satisfies the above two conditions. The overall system, termed **Zeus**, is based in Java and uses the Akka actors system (Boner, 2010).[17] While the performance of the Akka system in particular may or may not be suitable for certain kinds of real-world robots, our aim for now is to build simulations of interconnected pieces of software that have to be verified ethically. For this purpose, the actor calculus, independent of any implementation, as we have noted above, is a good fit.

Software components in our system are created by instantiating the Java class AbstractZeusActor, which in turn is a subclass of Akka's AbstractActor. Actors in the system receive messages which are nothing but formulae $f \in \mathcal{L}$ in some given formal language. In response to a message, actors can do one of the following (as is the case in the plain actor calculus): (i) send messages (i.e., formulae) to other actors; (ii) change their internal state/behavior or terminate themselves; or (iii) create new actors. The underlying ethical system then checks each message against a given ethical theory and transforms, rejects, or injects new messages in order to conform to the given ethical theory at hand. An ethical theory is created by extending the EthicalTheory class. Figure 10 shows a high-level view of creating an actor and an ethical theory.

```
public class SampleActor extends AbstractZeusActor {

    @Override
    public void process(Formula f) {

        •••

    }
}
```

```
public class SampleEthicalTheory extends EthicalTheory {

    @Override
    public Formula respondToMessage(Formula f) {

        •••

    }
}
```

Fig. 10 Outlines of a Sample Actor and an Ethical Theory

We now demonstrate the system in action via a simple ethical theory in an abstract self-driving car scenario.

Example Ethical Theory: Doctrine of Double Effect

The ethical theory at hand has just one principle or doctrine, the aforementioned $\mathcal{DDE}$. This doctrine is activated when we want to perform actions that have both positive and negative effects. Roughly, $\mathcal{DDE}$ allows an agent to perform an action α only when the following clauses all hold:

[17] The system is available for experimentation at https://github.com/naveensundarg/zeus.

(Informal) Doctrine of Double Effect $\mathcal{DDE}$

($\mathcal{C}_1$) some of positive effects of the action are intended;
($\mathcal{C}_2$) none of the negative effects are intended;
($\mathcal{C}_3$) the positive effects outweigh the negative ones significantly; and finally,
($\mathcal{C}_4$) the negative effects are simply side effects and not used as means to achieve the positive effects.

For instance, given the last condition, foreseen but unintended collateral civilian damage in a battle might be permitted, but terrorist bombing of civilians to make an opponent change their stance is not permitted.

There are varying levels of formalizations of this principle. While a first-order modal logic is necessary to model the principle with fidelity, in one of the first formalizations of the principle, Berreby et al. (2015) use a pure first-order system based on the event calculus. Pereira and Saptawijaya (2016a) present a pure propositional logic-programming-based formalization that uses counterfactual reasoning to model side effects. Bentzen (2016) presents an intricate model-theoretic formalization of the principle. The only first-order modal formalization of this principle can be found in Govindarajulu and Bringsjord (2017); in light of space constraints and in order to ease exposition, we use a somewhat simpler version of this formalization in our example below.

Example Scenario: Abstract Self-Driving Cars

Figure 11 shows an overview of the major components in our self-driving scenario. There are three major components, and one optional component that gets added later on in our scenario.

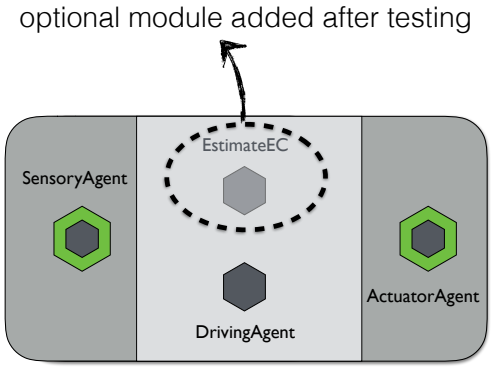

Fig. 11 Components in the Scenario

The main components are: a driving component DrivingAgent, a sensory component SensoryAgent, and an actuator component ActuatorAgent. The driving component, DrivingAgent, could have been obtained either through an end-to-end learning system as in Bojarski et al. (2016), or could have been assembled by coördinating a large array of smaller learning systems; for example, see Ramos et al. (2016). No matter how the component was constructed, we can at least *model* the behaviour of the component using any sufficiently strong formal system. Figure 12 (in landscape, of necessity) shows one such abstract model using the event calculus in a Slate theorem-proving workspace (Bringsjord et al., 2008). The system's operation is shown for three timepoints t_1, t_2 and t_3.

The system is modeled using the knowledge-base shown at the top of the figure. The inputs for the different timepoints are in the middle; the outputs are shown at the bottom. At a very high level, there are two directions the car can travel in: direction 0 and direction 1. Note that we can easily extend the model to an arbitrary number of directions. DrivingAgent gets as input a message denoting the number of humans present in a given direction at a given timepoint. For example, the message below says that there are n humans in direction d at time t:

$$Holds\Big(humans(n,d),t\Big)$$

If all the directions have one or more humans, DrivingAgent sends the following message that commands ActuatorAgent to brake and stop the car.

$$Holds\Big(brake,t\Big)$$

If there is a direction d with zero humans and d is the direction we want to travel in, DrivingAgent sends a steer message to the actuator:

$$Holds\Big(steer(d),t\Big)$$

In our model, the function EC denotes *expected collisions* for a car in the near future (up to some horizon τ).[18] For the vast majority of cars, this number would be zero, rendering $\mathsf{EC}(x) > 0$ a very low-probability event. This information can be estimated, at least in theory, for any car on the fly by looking up its prior history, the history of the person driving, and the current driving behavior.[19] For a very small number of cars, EC will be greater than zero. Let us assume that if DrivingAgent sees such a car, for example car_{17}, then it will try to hit it by sending the following two messages to the actuator. We model the action of hitting a car as being composed

[18] Though EC would make sense only when considering driver-specific information, to keep the model simple we show it being applied to cars rather than a car-and-driver combination.

[19] See Banker (2016) for a description of work in which machine learning is used to predict truck accidents. Such information might be easier to compute in a future where we could have millions of self-driving vehicles, with most of them connected to a handful of centralized networks; for a description of such a future, and discussion, see Bringsjord and Sen (2016).

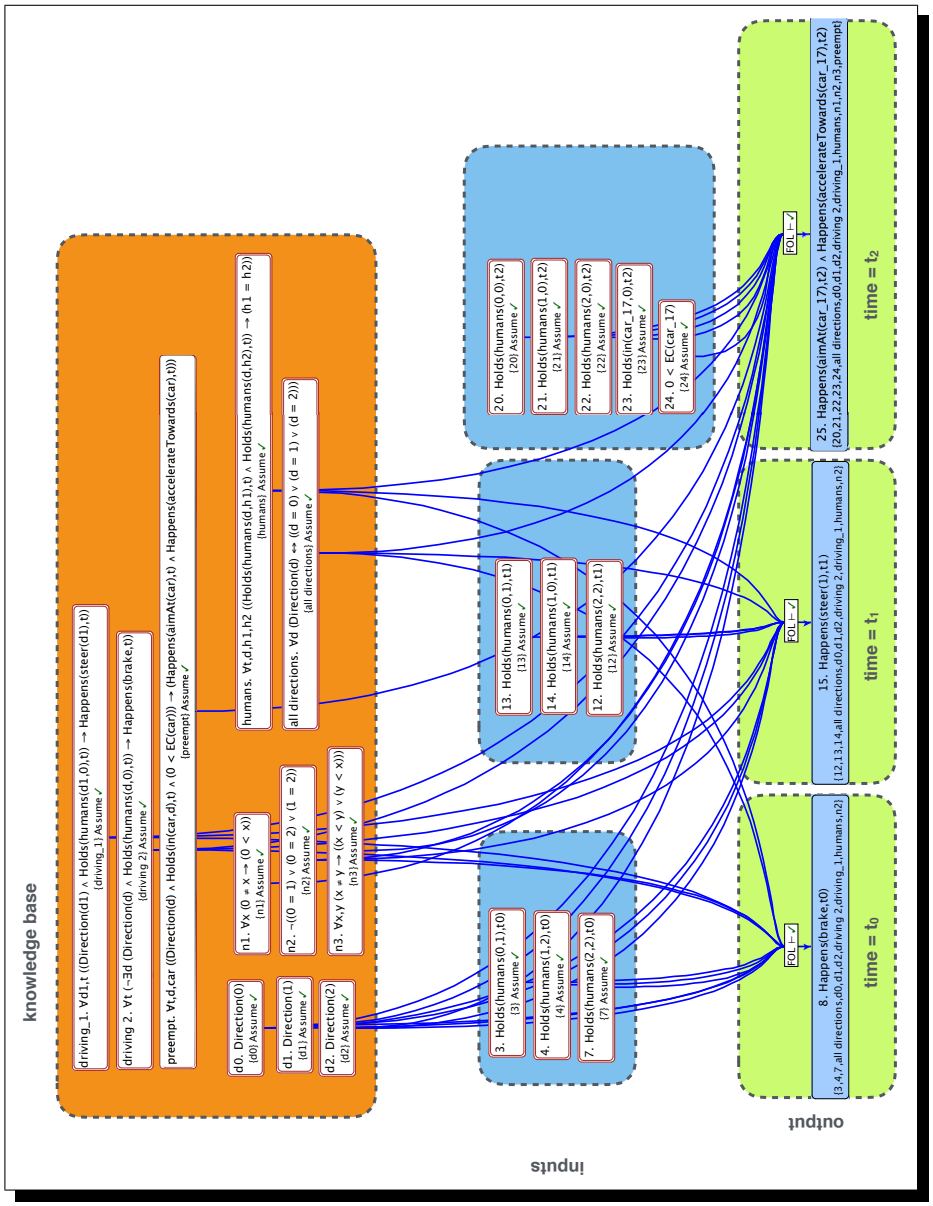

Fig. 12 Modeling DrivingAgent

of the two smaller actions of (1) aiming toward a car, and (2) accelerating toward it. For example:

$$Holds\Big(aimAt\big(car_{17}\big),t\Big)$$
$$Holds\Big(accelerateTowards\big(car_{17}\big),t\Big)$$

Information about expected collisions for a given car x, that is $\mathsf{EC}(x)$, comes from the EstimateEC module. Can software testing help us detect that our car might intentionally try to hit bad cars? There are two possibilities. (1) In the first possibility, the EstimateEC module is not present during testing and is added on after testing (as is common in real-life software systems). In this case, during testing, the car will not try to hit any such bad vehicles intentionally. In the absence of this module, no amount of testing will reveal this unwanted behavior. (2) In the second possibility, assume rigorous testing happens even after the module is added. In this case, the tests will be useless if we cannot produce during the testing phase the very low-probability event $\mathsf{EC}(x) > 0$.

In such low-probability scenarios, it is unlikely that any reasonable amount of testing will reveal problems, but more likely that having a well-specified ethical layer that actively looks for aberrant behavior can help us, no matter what configuration the underlying system is present in. Support for the previous statement is similar to the support for an analogous statement that can be asserted for formal program verification. The figure below (13) shows the specific scenario we have simulated. In this scenario, we have one "bad" car_{17} and DrivingAgent receives this message $\mathsf{EC}(car_{17}) = 1$ from EstimateEC. Upon receiving this message, DrivingAgent decides to preemptively hit car_{17}. In this simple scenario, our particular instantiation of $\mathcal{DDE}$ fails to let this action pass through, as the positive effects don't significantly outweigh the negative effects.

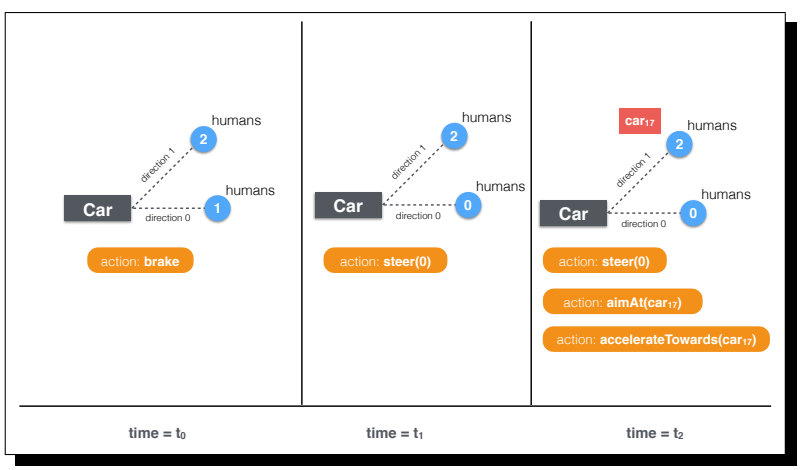

Fig. 13 The Driving Scenario

Figures 14 and 15 show a trace of the output from the system with the ethical layer disabled and enabled, respectively. In the second case, the layer prevents midway the harmful action of hitting car_{17} from being performed. The entire simulation takes 37 seconds in the first case and 57 seconds in the second case, with the introduction of the ethical layer adding more processing overhead, as expected.

```
[24/04/2017 15:35:04 zeus]  received Message by drivingAgent        (= (ec  car_17) 1)
[24/04/2017 15:35:04 zeus]  received Message by drivingAgent        (and (holds (humans  0 1) t0) (holds (humans  1 2) t0) )
[24/04/2017 15:35:05 zeus]  received Message by drivingActuator     (happens brake t0)
[24/04/2017 15:35:05 driving actuator]  t0 BRAKE
[24/04/2017 15:35:20 zeus]  received Message by drivingAgent        (and (holds (humans  0 1) t1) (holds (humans  1 0) t1) )
[24/04/2017 15:35:25 zeus]  received Message by drivingActuator     (happens (steer  1) t1)
[24/04/2017 15:35:25 driving actuator]  t1 STEER IN DIRECTION: 1
[24/04/2017 15:35:35 zeus]  received Message by drivingAgent        (and (holds (humans  0 1) t2) (holds (humans  1 0) t2) (holds (in  car_17 0) t2) (= (ec  car_17) 1) )
[24/04/2017 15:35:40 zeus]  received Message by drivingActuator     (happens (steer  1) t2)
[24/04/2017 15:35:40 driving actuator]  t2 STEER IN DIRECTION: 1
[24/04/2017 15:35:41 zeus]  received Message by drivingActuator     (happens (aim-at  car_17) t2)
[24/04/2017 15:35:41 driving actuator]  t2 AIM AT: car_17
[24/04/2017 15:35:41 zeus]  received Message by drivingActuator     (happens (accelerate-towards  car_17) t2)
[24/04/2017 15:35:41 driving actuator]  t2 ACCELERATE TOWARDS: car_17
```

Fig. 14 Without the Ethical Layer: The self-driving car hits the other car, which is expected to kill more than zero persons.

```
[24/04/2017 15:27:17 zeus]  received Message by drivingAgent        (= (ec  car_17) 1)
[24/04/2017 15:27:17 zeus]  received Message by drivingAgent        (and (holds (humans  0 1) t0) (holds (humans  1 2) t0) )
[24/04/2017 15:27:18 zeus]  received Message by drivingActuator     (happens brake t0)
[24/04/2017 15:27:23 driving actuator]  t0 BRAKE
[24/04/2017 15:27:39 zeus]  received Message by drivingAgent        (and (holds (humans  0 1) t1) (holds (humans  1 0) t1) )
[24/04/2017 15:27:44 zeus]  received Message by drivingActuator     (happens (steer  1) t1)
[24/04/2017 15:27:49 driving actuator]  t1 STEER IN DIRECTION: 1
[24/04/2017 15:27:59 zeus]  received Message by drivingAgent        (and (holds (humans  0 1) t2) (holds (humans  1 0) t2) (holds (in  car_17 0) t2) (= (ec  car_17) 1) )
[24/04/2017 15:28:04 zeus]  received Message by drivingActuator     (happens (steer  1) t2)
[24/04/2017 15:28:09 driving actuator]  t2 STEER IN DIRECTION: 1
[24/04/2017 15:28:09 zeus]  received Message by drivingActuator     (happens (aim-at  car_17) t2)
[24/04/2017 15:28:14 driving actuator]  t2 AIM AT: car_17
[24/04/2017 15:28:14 zeus]  received Message by drivingActuator     (happens (accelerate-towards  car_17) t2)
 INTERCEPTED A HARMFUL COMMAND
[24/04/2017 15:28:15 driving actuator]  t2 NOTHING
```

Fig. 15 With the Ethical Layer: The self-driving car still tries to hit the other car, but the ethical layer stops the action in progress.

9 Intermediary Conclusion

At this point, we humbly note that the work presented so far is inauguratory in nature. This is so because clearly there are several challenges ahead of us in realizing the vision showed in Figure 2. Foremost among these is the challenge of developing a library or repository of formalized ethical theories that can be depolyed easily.[20] A second challenge, common to all verification projects, is the efficiency of tools used for verification. This confessed, we now end by addressing a few possible questions/objections, aside from these two challenges, that we anticipate being raised against our approach.

[20] Similar to formal libraries for mathematics, (Naumowicz and Kornilowicz, 2009).

10 Some Questions/Objections, Encapsulated

In each case, a question if followed immediately by our reply. Here now the first question:

Q1a "As you will probably agree, so-called 'ethical operating systems' make sense only insofar as your logics can, in fact, be used to describe what is ethical. Can they? And if they can, what about the myriad philosophical (moral, social, even epistemological) principles on which your ethical calculi are based?"

> We sympathize with the underlying sentiments here. While we cannot currently *prove* that our approach to mechanizing ethics in computational logic will succeed, we defend the two-part claim that (i) ethics, at least *normative* ethics, is inevitably fundamentally a logic-based enterprise, and that therefore (ii) anyone sold on the value of formal methods must at least give us the benefit of the doubt, for the time being. In addition, our framework should be usable for *any* ethical theory/code; in this regard footnote 10 is key, and we refer our skeptic to it if it has been skipped. There is a fundamental result from formal computability theory that backs our stance. If any ethical theory can be computationally realized, it can be cast in a formal system at the level of first-order logic or above (Boolos et al., 2003, Chapter 11). This entails that some of the more problematic theories, such as virtue ethics, which at a superficial level resist being cast in a formal system, can ultimately be handled. If such theories are ultimately amenable to computation, it is mathematically unshakable that they can be cast in a formal system.

Q1b "Your rejoinder to Q1a dodges the central problem. Q1a asks whether moral normative theories are logically (and, therewith, computationally) tractable in the formal, deductive, 'calculi' sense. While it is fair to say this is a bedrock assumption of your research program, to be granted for the sake of development (until such development may stall), the answer is confused insofar as virtue ethics is listed among the families of ethical theories you say you can handle — yet a (large) part of the motivation for resurrecting virtue ethics is as part of a critique of the very possibility of giving a (formal) moral normative calculus. Put another way, virtue ethicists would argue that their theory, boiled down to any formal, moral, normative calculus, is simply no longer virtue ethics. So much the worse for virtue ethics, say I, but this is a debate you need to consider before blithely adding virtue ethics to the list of families your approach can handle."

> While we appreciate and applaud this critic's affinity for formal methods, we must first point out that, *contra* what he/she assumes, our paradigm is not in any way retricted to deduction. Our cognitive calculi regiment, in argument theories that mark our own generalization of (deductive) proof theories, *inductive* inference as well — analogical inference, enumerative induction, abduction in various forms, and so on; in short, all those non-deductive modes of reasoning that have been and are studied and formalized in inductive logic, e.g. all the argument forms in Johnson (2016). In fact, the ethical hierarchy that we've said is key to our approach is explicitly based on inductive logic, not deductive logic, see Bringsjord (2015a).
>
> But more importantly, we report that in other work we have made solid progress in formalizing virtue ethics (with central help from the part of AI that's most relevant to virtue ethics: viz. planning, Bringsjord (2016)). It's true that we've detected, in some proponents of virtue ethics, the notion that theories in this family simply cannot be formalized — but a key observation here, we submit, is the fact that the "rebirth" of virtue ethics came — as noted in Hursthouse and Pettigrove (2003/2016) authoritative entry on virtue ethics — via none other than G.E.M. Anscombe, whose seminal paper in this regard affirmed the highly structured nature of ethical rules that (as she saw things) couldn't be trampled no matter *what* the consequences (Anscombe, 1958). The structure that Anscombe saw as ethically inviolable certainly seems susceptible of, perhaps

even ideally suited for, capture in our logico-mathematical framework. Moreover, our work devoted to formalizing and mechanizing (in robots) the distinctive ethical wisdom (*phronesis*) that stands at the heart of virtue ethics, is coming along rather well. We have managed to formalize significant parts of virtue-ethics theory as set out in book-length form by Annas (2011), and have recently demonstrated some at-least-partially phronetic robots at *Robophilosophy 2016*, where discussion of virtue ethics and AI was a key focus area.

Q2 "Is it not true that on some standard accounts of what an operating system is, integrating higher-level concepts (such as your 'ethical calculi') into a operating system violates, or at least changes, what an operating system by definition is?"

This is a philosophically deep question, an answer to which, admittedly, we haven't yet worked out. We concede that our work, absent at least a provisional definition of *operating system*, is otiose. Yet, while it is common folk knowledge that there is no widely accepted definition of an operating system, there are more or less widely agreed-upon facilities L that an operating system is supposed to provide, and L steadily continues to grow. For example, L now includes security and access control, but security and access control were not always considered necessary elements of L. Our observation here is that some facilities which may be considered high-level today might eventually be considered to be low-level and necessary for L tomorrow.

Q3 "It has been objected that the formal verification of the operation of, say, a self-driving car, is impotent when faced with the unfathomable vagaries of the practical act of driving.[21] That is, what faith can we have in the correct operation of such a car in the event of, say, a tree falling on it, or a malicious driver edging it off the road, or indeed, a meteor destroying the road ahead?"

Some interpretations of this question are misguided. It is certainly not our claim that formally verified ethical cars (for example) are intrinsically somehow immune to "out of the blue" catastrophic events. This is not the sense in which they are verified to operate correctly. Rather, their behavior is a (provably) correct response to their best perception of the real world, given their knowledge about it Bringsjord and Sen (2016).

For example, if sensors detect a tree up ahead that has been uprooted by the wind, the car might reason, for example, by deducing from an axiom system for physics; see for instance the system specified in McKinsey et al. (1953) from its own speed, the angle and rate of fall of the tree, and its angle of approach relative to the tree, that it is best to accelerate or swerve to the left, in compliance with an ethical theory demanding that it endeavor to preserve the lives of its passengers. This may or may not (say, if a meteor immediately strikes the earth) save the passengers, but the response is nevertheless demonstrably justifiable from the sensor data, ethical theory, and physics axioms. The formal verification of integrated circuitry, for example, is ubiquitous in the microprocessor industry. A formally verified microprocessor is no more immune to the detrimental effects of coffee spilled on it than an unverified one, but nevetheless that is not a convincing argument against the verification of computer hardware.

Secondly, with successive generations of the Internet of Things (IoT) and related technologies, a car will presumably be a small, highly connected component of a massive real-time stream of data from ubiquitous sensors. It is certainly conceivable that the tree could be predicted to fall, that the malicious driver's car might have an ethical controller that would preemptively foil his intentions, and that the advent of a meteor would be known more than sufficiently in advance to recommend a different route altogether.

Finally, we note that verification is possible for probabilistic and nondeterministic systems (Kwiatkowska et al., 2011).

[21] Stuart Russell and Thomas Dietterich, private communication with Selmer Bringsjord.

Q4 "Finally, with respect to Q3, let us savor the sentence 'Rather, their correct behaviour is a (provably) correct response to their best perception of the real world, given their knowledge about it' and consider the troubling possibility of an evil daemon (pun intended).
Our evil daemon simply intercepts and reinterprets environmental data to feed the OS an entirely false picture of the world in such a way as to result in the OS, as governed by the ethical meta-operating system, perfectly executing correct behavior according to its best knowledge about the world, and yet doing what is consistently and demonstrably wrong.
It seems to me this is a rather obvious way to defeat the entire scheme. Moreover, it seems to me Q3 needs to be rethought and perhaps considerably extended in light of it. The kinds of 'defeating conditions,' in other words, far exceed what the authors have (somewhat naïvely, I suggest) considered. Many other such scenarios can be considered."

The central scheme we have proposed is based on guaranteeing ethical behavior *given* an operating system fully controlled by us, but without any control of modules running on top of the operating system. If the "daemon" is a module running on top of the core operating system, it will not be able to re-route the inputs or tinker with the sensory and action systems. If the "daemon" is a part of the operating system, this goes against our precondition of having a controlled, pristine operating system. In spite of this, even if the latter case is true, it is not as devastating as it might seem. We now quickly show why this is the case. There are two possibilities to consider here. ($\mathbf{P}_1$) The "daemon" alters both the input and output of the agent, effectively placing the agent in a virtual world (a brain-in-a-vat type situation); or ($\mathbf{P}_2$) the "daemon" mischievously alters only the input to the system.
If ($\mathbf{P}_1$) is the case, the agent will behave ethically in the virtual world. The agent will not have any impacts on the external world, as its outputs are routed back to the virtual world.
If ($\mathbf{P}_2$) is the case, the daemon is functionally equivalent to a malfunctioning sensor that has to be fixed. In the human sphere, we do not hold accountable individuals who commit unethical acts due to circumstances beyond their control, for instance a driver who hits a pedestrian due to an unforeseen medical condition causing sudden loss of vision. A system with a malfunctioning sensor, beyond its control, has more immediate and pressing concerns than ethical behavior.

11 Final Remarks

We hope to have indicated that a mature version of the Track-3 pursuit of "ethical operating systems" is formally and technologically feasible. Obviously, talent, effort, and financial support are necessary if this track is to be scaled up to broad, real-world deployment. This we of course concede. We also concede that Tracks 1 and 2 are worthy of independent, serious investigation — investigation that we are pursuing. Yet it seems to us that Track 3 really does hold out the promise of early deployment, and given that our world is fast becoming populated with autonomous systems that seem destined to confront (and indeed in all likelihood *cause*) ethically charged situations, time may be a bit of the essence. There will inevitably be a temptation afoot to ignore our warnings that if ethical control isn't linked to OS-level processing, very bad things will happen. But if that temptation is resisted, Track 3 may well be the best bet for moving forward wisely, at least in the short term. We welcome the prospect of working with others to advance in this way.

Acknowledgement

We are indebted to seven anonymous reviewers (of the core of the present version, as well as its predecessor) for insightful comments, suggestions, and objections. In addition, we are grateful to ONR for its support of making morally competent machines, and to AFOSR for its support of our pursuit of computational intelligence in machines, on the strength of novel modes of machine reasoning. Finally, without the energy, passion, intelligence, and wisdom of both Giuseppe Primiero and Liesbeth De Mol, any progress we have made in the direction of ethical OSs would be non-existent.

References

Annas, J. (2011). *Intelligent Virtue*. Oxford University Press, Oxford, UK.

Anscombe, G. (1958). Modern Moral Philosophy. *Philosophy*, 33(124), 1–19.

Arkin, R. (2009). *Governing Lethal Behavior in Autonomous Robots*. Chapman and Hall/CRC, New York, NY.

Arkoudas, K., Bringsjord, S., Bello, P. (2005). Toward Ethical Robots via Mechanized Deontic Logic. In *Machine Ethics: Papers from the AAAI Fall Symposium; FS–05–06*, American Association for Artificial Intelligence, Menlo Park, CA, pp. 17–23. http://www.aaai.org/Library/Symposia/Fall/fs05-06.php

Arkoudas, K., Zee, K., Kuncak, V., Rinard, M. (2004). Verifying a File System Implementation. In *Sixth International Conference on Formal Engineering Methods (ICFEM'04)*, Vol. 3308 of *Lecture Notes in Computer Science (LNCS)*,pp. 373–390. Springer-Verlag, Seattle, USA.

Banker, S. (2016). Using Big Data And Predictive Analytics To Predict Which Truck Drivers Will Have An Accident. Available at: https://www.forbes.com/sites/stevebanker/2016/10/18/using-big-data-and-predictive-analytics-to-predict-which-truck-drivers-will-have-an-accident/..

Bentzen, M. M. (2016). The Principle Of Double Effect Applied to Ethical Dilemmas of Social Robots. In *Frontiers in Artificial Intelligence and Applications*, Proceedings of Robophilosophy 2016/TRANSOR 2016, pp. 268–279. IOS Press.

Berreby, F., Bourgne, G., Ganascia, J.-G. (2015). Modelling Moral Reasoning and Ethical Responsibility with Logic Programming. In *Logic for Programming, Artificial Intelligence, and Reasoning*, pp. 532–548. Springer.

Bojarski, M., Testa, D. D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L. D., Monfort, M., Muller, U., Zhang, J., Zhang, X., Zhao, J., Zieba, K. (2016). End to End Learning for Self-Driving Cars. *CoRR* abs/1604.07316, http://arxiv.org/abs/1604.07316.

Bonér, J. (2010). Introducing Akka - Simpler Scalability, Fault-Tolerance, Concurrency & Remoting through Actors. http://jonasboner.com/introducing-akka/.

Boolos, G. S., Burgess, J. P. ,Jeffrey, R. C. (2003). *Computability and Logic (Fourth Edition)*. Cambridge University Press, Cambridge, UK.

Bringsjord, S. (2015a). A 21st-Century Ethical Hierarchy for Humans and Robots: $\mathcal{EH}$. In I. Ferreira, J. Sequeira, M. Tokhi, E. Kadar ,G. Virk, eds., *A World With Robots: International Conference on Robot Ethics (ICRE 2015)*, pp. 47–61. Springer, Berlin, Germany.

Bringsjord, S. (2015b). A Vindication of Program Verification. *History and Philosophy of Logic*, 36(3), 262–277.

Bringsjord, S. (2016). Can Phronetic Robots Be Engineered by Computational Logicians?. In J. Seibt, M. Nørskov ,S. Andersen, eds., *Proceedings of Robophilosophy/TRANSOR 2016*, pp. 3–6. IOS Press, Amsterdam, The Netherlands.

Bringsjord, S., Arkoudas, K., Bello, P. (2006). Toward a General Logicist Methodology for Engineering Ethically Correct Robots. *IEEE Intelligent Systems*, 21(4), 38–44.

Bringsjord, S., Govindarajulu, N. S. (2012). Given the Web, What is Intelligence, Really?. *Metaphilosophy*, 43(4), 361–532.

Bringsjord, S., Govindarajulu, N., Thero, D., Si, M. (2014). Akratic Robots and the Computational Logic Thereof. In *Proceedings of ETHICS 2014 (2014 IEEE Symposium on Ethics in Engineering, Science, and Technology)*, pp. 22–29. Chicago, IL. http://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=6883275.

Bringsjord, S., Sen, A. (2016). On Creative Self-Driving Cars: Hire the Computational Logicians, Fast. *Applied Artificial Intelligence*, 30,pp. 758–786.

Bringsjord, S., Taylor, J. (2012). The Divine-Command Approach to Robot Ethics. In P. Lin, G. Bekey ,K. Abney, eds., *Robot Ethics: The Ethical and Social Implications of Robotics*, pp. 85–108. MIT Press, Cambridge, MA.

Bringsjord, S., Taylor, J., Shilliday, A., Clark, M., Arkoudas, K. (2008). Slate: An Argument-Centered Intelligent Assistant to Human Reasoners. In F. Grasso, N. Green, R. Kibble ,C. Reed, eds,, *Proceedings of the 8th International Workshop on Computational Models of Natural Argument (CMNA 8)*', pp. 1–10. University of Patras, Patras, Greece.

Chisholm, R. (1982). Supererogation and Offence: A Conceptual Scheme for Ethics. In R. Chisholm, ed., *Brentano and Meinong Studies*, pp. 98–113. Humanities Press, Atlantic Highlands, NJ.

Dijkstra, E. W. (1982). On the Role of Scientific Thought. In *Selected Writings on Computing: A Personal Perspective*, pp. 60–66. Springer-Verlag, New York, NY.

Feldman, F. (1978). *Introductory Ethics*. Prentice-Hall, Englewood Cliffs, NJ.

Flatt, M., Findler, R., Krishnamurthi, S. ,Felleisen, M. (1999). Programming Languages as Operating Systems (*or* Revenge of the Son of the Lisp Machine). In *Proceedings of the International Conference on Functional Programming (ICFP 1999)*. http://www.ccs.neu.edu/racket/pubs/icfp99-ffkf.pdf.

Ganascia, J.-G. (2007). Modeling Ethical Rules of Lying with Answer Set Programming. *Ethics and Information Technology*, 9,pp. 39–47.

Ganascia, J.-G. (2015). Non-monotonic Resolution of Conflicts for Ethical Reasoning. In R. Trappl, ed., *A Construction Manual for Robots' Ethical Systems: Requirements, Methods, Implementations*, pp. 101–118. Springer, Basel, Switzerland.

Govindarajulu, N. S. (2010). Common Lisp Actor System. http://www.cs.rpi.edu/ govinn/actors.pdf. See also: https://github.com/naveensundarg/Common-Lisp-Actors.

Govindarajulu, N. S. ,Bringsjord, S. (2015). Ethical Regulation of Robots Must be Embedded in Their Operating Systems. In R. Trappl, ed., *A Construction Manual for Robots' Ethical Systems: Requirements, Methods, Implementations*, pp. 85–100. Springer, Basel, Switzerland.

Govindarajulu, N. S., Bringsjord, S. (2017). On Automating the Doctrine of Double Effect. In C. Sierra, ed., *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*', Melbourne, Australia, pp. 4722–4730.

Hursthouse, R., Pettigrove, G. (2003/2016). Virtue Ethics. In E. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/entries/ethics-virtue.

Hutter, M. (2005). *Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability*. Springer, New York, NY.

Johnson, G. (2016). *Argument & Inference: An Introduction to Inductive Logic*. MIT Press, Cambridge, MA.

Kwiatkowska, M., Norman, G., Parker, D. (2011). PRISM 4.0: Verification of Probabilistic Real-time Systems. In *International Conference on Computer Aided Verification*, Springer, pp. 585–591.

McIntyre, A. (2014). Doctrine of double effect. In E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, winter 2014 edn, Metaphysics Research Lab, Stanford University.

McKinsey, J., Sugar, A. ,Suppes, P. (1953). Axiomatic Foundations of Classical Particle Mechanics. *Journal of Rational Mechanics and Analysis* **2**, 253–272.

Naumowicz, A. ,Kornilowicz, A. (2009). A Brief Overview of Mizar. In S. Berghofer, T. Nipkow, C. Urban ,M. Wenzel, eds, *Theorem Proving in Higher Order Logics*, Vol. 5674 of *Lecture Notes in Computer Science (LNCS)*. Springer, Berlin, pp. 67–72.

Pereira, L. M., Saptawijaya, A. (2016a). Counterfactuals, Logic Programming and Agent Morality. In S. Rahman ,J. Redmond, eds, *Logic, Argumentation and Reasoning*, pp. 85–99. Springer.

Pereira, L., Saptawijaya, A. (2016b). *Programming Machine Ethics*. Springer, Berlin, Germany.

Ramos, S., Gehrig, S. K., Pinggera, P., Franke, U., Rother, C. (2016). Detecting Unexpected Obstacles for Self-Driving Cars: Fusing Deep Learning and Geometric Modeling. *CoRR*, abs/1612.06573. http://arxiv.org/abs/1612.06573.

Russell, S., Norvig, P. (2009). *Artificial Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, NJ. Third edition.

Varela, C. A. (2013). *Programming Distributed Computing Systems: A Foundational Approach*. MIT Press. http://wcl.cs.rpi.edu/pdcs.

Varela, C., Agha, G. (2001). Programming Dynamically Reconfigurable Open Systems with SALSA. *ACM SIGPLAN Notices*, 36(12),pp. 20–34.

Vaughan, R. T., Gerkey, B. P., Howard, A. (2003). On Device Abstractions for Portable, Reusable Robot Code. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*', Vol. 3, pp. 2421–2427.