IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received April 3, 2019, accepted May 2, 2019, date of publication May 10, 2019, date of current version May 23, 2019.

*Digital Object Identifier 10.1109/ACCESS.2019.2916203*

# A Software Exoskeleton to Protect and Support Citizen's Ethics and Privacy in the Digital World

**MARCO AUTILI[1], DAVIDE DI RUSCIO[1], PAOLA INVERARDI[1], PATRIZIO PELLICCIONE[1,2], AND MASSIMO TIVOLI[1]**

[1]DISIM, University of L'Aquila, 67100 L'Aquila, Italy
[2]Department of Computer Science and Engineering, Chalmers University of Technology | University of Gothenburg, SE-405 30 Gothenburg, Sweden

Corresponding author: Patrizio Pelliccione (patrizio.pelliccione@gmail.com)

**ABSTRACT** Citizens of the digital world are threatened. The digital systems that surround them are increasingly able to make autonomous decisions over and above them and on their behalf. They feel that their moral rights, as well as the social, economic, and political spheres, can be affected by the behavior of such systems. Although unavoidable, the digital world is becoming uncomfortable and potentially hostile to its users as human beings and as citizens. Notwithstanding the introduction of the GDPR and of initiatives to establish criteria on software transparency and accountability, users feel vulnerable and unprotected. In this paper, we present EXOSOUL, an overarching research framework that aims at building a software a personalized exoskeleton that enhances and protects users by mediating their interactions with the digital world according to their own ethics of actions and privacy of data. The exoskeleton disallows or adapts the interactions that would result in unacceptable or morally wrong behaviors according to the ethics and privacy preferences of the users. With their software shield, users will feel empowered and in control, and more in the balance of forces with the other actors of the digital world. To reach the breakthrough result of automatically building a personalized exoskeleton, EXOSOUL identifies multidisciplinary challenges never touched before: 1) defining the scope for and inferring citizen's ethical preferences; 2) treating privacy as an ethical dimension managed through the disruptive notion of active data; and 3) automatically synthesizing ethical actuators, i.e., connector components that mediate the interaction between the user and the digital world to enforce her ethical preferences. In this paper, we discuss the research challenges of EXOSOUL in terms of their feasibility and risks.

**INDEX TERMS** Autonomous systems, AI, artificial intelligence, privacy, ethics, software engineering.

## I. INTRODUCTION

In their ordinary life, citizens in the digital world continuously interact with software systems, e.g., by using a mobile device or from on board of a (autonomous) car. These systems are increasingly autonomous in making decisions over and above the users or on behalf of them [38]. Often, their autonomy exceeds the system boundaries and invades user prerogatives. As a consequence, ethical issues – privacy ones included (e.g., unauthorized disclosure and mining of personal data, access to restricted resources) – are emerging as matters of utmost concern since they impact on the moral rights of each human being and affect the social, economic, and political spheres[1] [8], [23], [29]. Besides the philosophical aspects, the way to approach these problems is twofold: *regulatory* and *technical*. Europe has recently introduced the GDPR legislation for data protection [6], it is at the forefront on the regulation of autonomous vehicles [10], [12], while a common EU approach to liability rules and insurance for connected and autonomous vehicles is under discussion [32]. Besides scientific community and big companies, Europe is also proposing initiatives to identify problems and establish criteria to develop algorithms and systems that embed autonomous capabilities [4], [7], [9], [14], [31]. As a matter of fact, the digital world is being recognized as potentially hostile to citizens. The initiatives proposed so far go in the

The associate editor coordinating the review of this manuscript and approving it for publication was Fabrizio Messina.

[1]https://www.politico.eu/article/cambridge-analytica-facebook-data-brittney-kaiser-privacy/

direction to make the world less hostile by introducing new laws, from the regulatory side, and transparency and accountability criteria in software development, from the technical side. Regulation is important as well as in-depth insights into the technology. However, we are fully aware that achieving full adherence to regulation and transparency criteria is very difficult or even impossible in practice. We are facing a paradox: human beings are recognized as central actors, *the sensitive targets*; but they are passive consumers in the digital world, and the power and the burden to preserve their rights remain in the hands of the (software-) systems producers. In the *mangrove societies* – Floridi's powerful metaphor[2] – human beings are unprotected in their interactions with the digital world. The great challenge, unattempted so far, is to comprehensively empower them.

As discussed in [38], "*there is the need to rethink the role of the various actors in the digital world by empowering the users of the digital technology both when they operate as citizens and as individuals*". In this paper we discuss EXOSOUL, an overarching research framework that aims at equipping humans with an **automatically generated exoskeleton**, i.e. a software shield that protects them and their personal data via the mediation of all interactions with the digital world that would result in unacceptable or morally wrong behaviors according to their ethical and privacy preferences. The exoskeleton can take a whole spectrum of forms: from customized soft-libraries that the individual may deploy on the machines being used, to a sophisticated software interface that an individual may "wear", eventually deployed on a body chip. **Empowering the users with a personalized exoskeleton will introduce more symmetry of power in the present digital world and will effectively put humans in the center**. Exoskeletons development also opens unprecedented business opportunities in the field of societal-friendly applications as happened in the case of open source software, which promoted the principles of free software against the monopoly proprietary software producers [11]. Furthermore, bringing back to the user part of the (digital) control helps to solve liability issues in autonomous systems by readdressing responsibility to users according to their specified ethics. In fact, by means of software exoskeletons on the one hand users will be protected by mediating their interactions with the digital world, and on the other hand they will be responsible of the consequences of theirs (ethical) decisions.

**Paper roadmap –** The paper is organized as follows. Section II discusses the main challenges to be faced to enable the automatic synthesis of software exoskeletons out of user's ethics and privacy preferences. Section III introduces a reference use case in the automotive domain. Section IV overviews the state of the art technologies our research framework, as described in Section V, bases on. Finally, Section VI concludes the paper.

[2]In the digital world it is impossible to distinguish whether we are online or offline - instead we are *onlife*, as it is impossible to understand whether the water in the estuary - where the river meets the sea - is sweet or salty [33].

## II. THE CHALLENGES AHEAD

We address the challenge of automatically synthesizing a software exoskeleton starting from the ethics and privacy preferences of the user. In the ethical sphere, this requires to answer several cutting edge research questions concerning the need to:

- Identify a space of ethics and privacy preferences for users, to assess their compatibility with regulations, and to orchestrate interactions of users endorsing different preferences, so as to prevent deadlocks and to promote best ethical practices in digital societies;
- Infer ethics and privacy preferences from the user. Note that, this is a very challenging task given that neither a person nor a society apply moral categories separately, rather everyday morality is in constant flux among norms, utilitarian assessment of consequences, and evaluation of virtues.

We define the exoskeleton by considering two specific classes of interactions that citizens have with the digital world. The first one concerns interactions that involve the exchange of personal data, and that as such impact the privacy dimension, notably interactions with mobile apps through mobile devices. Until now, data are considered as passive entities and the logic implementing their life-cycle is decoupled from the data itself. For each datum that is shared over the Internet, the owner loses its track and control [41]. Such problems have been mitigated by means of regulatory (e.g., GDPR) and technical attempts. Unfortunately, these attempts solve the mentioned issues only partially, e.g. how data being modeled, how to enforce privacy concerns, or propose access control policies. We propose a disruptive approach that changes the passive nature of data by introducing **active data**. As part of the exoskeleton, active data encapsulate data with mechanisms that govern their creation, destruction, use, and sharing according to the owner ethical preferences. Destruction is the basic means to provide the *right to be forgotten*, which requires to equip data with an apoptosis mechanism[3] – synthesized from the user's ethical and privacy preferences – whose enactment depends on the use the digital world makes of the data, beyond parameters like time. The second one concerns the interaction with systems that are equipped with some degree of autonomy and that a user may want to ethically control to some extent. Autonomous vehicles and the so-called *trolley problem* represent a well-known limiting case, but other more ordinary cases exist [25]. As part of the exoskeleton, we will address the challenge of synthesizing, out of the user's ethical preferences, an **ethical actuator** able to intercept the interactions between the autonomous engine and the machine actuators and to prevent behaviors that are not admissible by the ethical preferences. Since this approach cannot be independent from the software the citizens are interacting with, by-product results of EXOSOUL will be requirements on the way the

[3]In biological terms, apoptosis, also called programmed cell death, is a mechanism that allows cells to self-destruct when stimulated by the appropriate trigger, internal or external to the cell - ref. Encyclopædia Britannica.

digital world needs to conform in order to interact with exoskeletons. This is as important as developing the shield since it establishes architecture and protocol requirements the systems producers need to comply with. EXOSOUL citizens will interact only with the part of the digital world that accepts their requirements. This breaks the monopoly of producers by introducing symmetry in the producer/user roles and new economic drivers in the digital market. At the same time, producers can be relieved from the liability burden by readdressing responsibility to users.

Summarizing, the high-level objective of EXOSOUL is to build a software exoskeleton that enhances and protects humans by mediating their interactions with the digital world according to their ethics of actions and privacy of data. From the technological point of view the major challenges that are emerging in order to ideate and develop innovative theories, methodologies and tools for achieving the vision of EXOSOUL are:

- *ch1* – is to conceive logic theories and supporting techniques to specify and infer user's ethical and privacy preferences.
- *ch2* – is to design the exoskeleton with its constituent active data and ethical actuator components, and to define newfangled techniques and tools for managing the exoskeleton life-cycle.
- *ch3* – is to define innovative synthesis techniques to generate personalized software exoskeletons.

These challenges will be investigated in Sections V-A, V-B, and V-C, respectively.

We plan to assess research and innovation outcomes of EXOSOUL by experimenting in the automotive domain. A further aim is to identify practical guidelines for software producers to support the deployment, execution, and run-time management of exoskeleton components.

As long term and social impact, we foreseen that EXOSOUL will revolutionize the de-facto standard processes of ethics management and personal data control. This can be done only by directly involving users. EXOSOUL contributes to societal impact by empowering citizens. The well founded ambition is to reach a wide adoption of exoskeletons from citizens so that big players (such as Android, Apple, and automotive OEMs) will be induced to change the way they manage control policies, users data, and ethics.

## III. USE CASE IN THE AUTOMOTIVE DOMAIN

(i) *Setting*: a parking lot in a big mall; (ii) *Resource contention*: two autonomous connected vehicles (named A and B hereafter), with one passenger each, are competing for the same parking lot. Passenger of vehicle A is pregnant. (iii) *Context*: A and B are rented vehicles, therefore, they are multi-user and have a default ethics that determines their decisions. The default ethics of A and B are utilitarian. Thus, the cars will look for the free parking lot that is closer to the point of interest, in case of contention the closest car gets in. The personal ethics of the passengers are transferred to vehicles from their mobile phones. (iv) *Action*: A and B are approaching the parking lot. B is closer, therefore it would take the parking lot. However, by communicating with A, it receives the information that the passenger in A is pregnant. Indeed, the exoskeleton of the passenger in A has disclosed such personal information through an active data that has a rule specifying that the data shall exist only within the parking lot. The exoskeleton of the passenger in B enacts her own ethics, which is a virtue ethic [3], and, consequently, actions are taken to leave the parking lot to A. Upon exiting the parking area, all the instances of the data regarding passenger A query the GPS and activate their own destruction. Finally, passenger B receives a notification about the ethical preferences that triggered the behavior of vehicle B and a positive feedback for the action she performed (this might encourage engagement with positive behaviors, similar to gamification).

This use case shows how personal privacy is strictly connected to ethics: by disclosing a personal information like this, the pregnant woman follows a utilitarian view which is related to the expectation that surrounding drivers might have a virtue personal ethic.

## IV. STATE OF THE ART

**Ethics and privacy –** As outlined in [38], the European Data Protection Supervisor (EDPS) in his strategy 2015-2019 established the Ethics Advisory Group with the mandate to reflect on the ethical implications emerging from the digital world. In [30], the EDPS defines the fundamental right to privacy and the protection of personal data that shall be guaranteed in order to preserve human dignity. It also calls for a 'big data protection ecosystem' that shall involve developers, businesses, regulators and individuals in order to provide 'future-oriented regulation', 'accountable controllers', 'privacy-conscious engineering', and 'empowered individuals'. EXOSOUL sets its actions on the last two elements of this ecosystem. In his 2018 report [29], the EDPS Ethical Advisory Board has provided a wide set of reflections on the notion of digital ethics that address the ''fundamental questions about what it means to make claims about ethics and human conduct in the digital age, when the baseline conditions of humanness are under the pressure of inter-connectivity, algorithmic decision-making, machine-learning, digital surveillance and the enormous collection of personal data …''. EXOSOUL shares this broader view of digital ethics and takes the challenge of re-conducting privacy concerns under more general ethical principles. We rely on the notion of digital ethics as presented in [33] as the branch of ethics that aims at formulating and supporting morally good solutions through the study of moral problems relating to personal data, (AI) algorithms and corresponding practices and infrastructures. A component of digital ethics is *hard* ethics that is defined and enforced by digital legislation [33], e.g., GDPR regulation [6]. However, legislation does not cover everything, nor should it. In the space not regulated by legislation the various entities of the digital world e.g., companies and citizens, should identify their role according to their digital ethics. This is the domain of *soft* ethics, which deals with

moral decisions over and above the existing regulation, without trying to by-pass or change the hard ethics. Soft ethics is exactly what EXOSOUL aims to support in the personal exoskeleton while we expect hard ethics to be implemented by the machine producers.

**Enabling technologies –** The EXOSOUL approach is based on two main software enabling technologies: specification techniques to specify ethics and privacy preferences, and architectural connector synthesis techniques to enforce interaction policies among distributed components forming a system. Software architectures represent another key competence that will be exploited in three different directions: (i) for architecting the exoskeleton, (ii) for enabling automated synthesis and enforcing techniques, and (iii) for developing recommendations in order to provide stakeholders with practical guidelines to enable the deployment, execution, and run-time management of the exoskeleton. The group has a consolidated experience in architectural languages [42], [43], architecture analysis [26], [47], architectural connectors synthesis [16]–[19], [39], [55], as well as architectural works in the robotic [34] or automotive domain [48], [58].

*Techniques to characterize privacy and ethics*: Privacy concerns data and has been historically addressed by means of permission systems that comprise both specification of access policies and their enforcement [49]. Recent works have addressed the problem of empowering the user in their interaction with mobile devices that are the most diffuse ways to access the digital world, e.g., [2], [50], [57]. In the Android context, which represents more than the 85% of the worldwide smartphone volume [1], this is done by asking the user to enter in the loop. In the new Android permission model [51], sophisticated and accurate user privacy-preference profiles are built via data collection and machine learning [57]; in our own project [50], the user is provided with a notation to specify and customize fine-grained permission levels according to her own subjective privacy concerns.

As far as ethics is concerned, the dichotomy between user ethical principles and autonomous systems insistently emerged in the autonomous cars domain, e.g., [25], [35]. In implementing ethics, we can distinguish: (i) approaches that try to cast general ethical categories (e.g., deontology and consequentialism) into elements of a mathematical problem (e.g., constraints and costs in a optimization problem) [44], [54], and (ii) iterative approaches for inferring and building less or more sophisticated ethical models out of ethical behaviors or decisions of the system [53].

*Architectural connector synthesis techniques for enforcing interaction policies*: Synthesis and enforcement of interaction policies have been addressed in different application domains, notably in the service-oriented domain as a solution for the realizability of choreographies. Our approach [16]–[19], [39], [55] is defined at the architectural level by synthesizing connectors behavioral models, and corresponding code, which coordinate the interactions among system's components in order to prevent mismatches, i.e., global system behaviors that violate the specified interaction properties. The synthesized connector model can be specialized to behave as a centralized coordinator [55] or a distributed one [19], a mediator [16], [39], or a distributed enforcer [17], [18], depending on the coordination issues, the protocol mismatches to solve, and the nature of the system.

Summarizing, we believe that theoretical foundations and tool support for specification and inference techniques, synthesis and enforcing techniques, as well as privacy and ethics awareness have reached the maturity to enable the creation of software personalized exoskeletons.

We will leverage both declarative specifications and experimental approaches driven by user data in order to produce personalized software ethics models and personalized set of active data. Synthesis and model-driven techniques will be used to derive behavioral models of the exoskeleton and the related code for the two application domains of interest. The synthesis process will be engineered so to generate exoskeletons that are robust and secure by exploiting state of the art solutions.

## V. EXOSOUL

In this section, we present the research themes that realize our research framework in accordance with Figure 1.

Research Theme 1 (RT1), which is detailed in Section V-A, investigates logic theories and supporting techniques for enabling users to infer and specify their ethical and privacy preferences. ''Privacy (P) and Digital Ethics (DE) Principles and Guidelines'' refer to *hard* privacy and *hard* ethics that are defined and enforced by digital legislation. Instead, ''Instruments to Specify and Infer P and DE User Preferences'' enable the user to define her *soft* privacy and *soft* ethics. We recall that this deals with what ought and ought not to be done over and above the existing regulation, without trying to by-pass or change the hard privacy and ethics.

As explained in RT1 below, the ''User-defined Domain independent P and DE Preferences'' will be defined via a top-down approach (<<define>> arrow), then refined and tuned up via a bottom-up approach (<<refine>> arrow from ''Demonstrators''). Research Theme 2 (RT2), which is detailed in Section V-B, conceptually defines the exoskeleton together with the techniques for manipulating it. Research Theme 3 (RT3), which is detailed in Section V-C, investigates innovative synthesis techniques to generate exoskeletons so to reflect user privacy and ethical preferences. In Figure 1, we highlight in light-gray the artifacts that are provided as input by the users and the domain experts. More precisely, user-defined domain-independent P and DE preferences are inferred and specified by end users of EXOSOUL through the instruments produced by RT1. In turn, this artifact is the input of the domain-independent exoskeleton synthesis. The other artifacts that are highlighted in light-gray (see the Domain-dependent Specifications box) are provided by domain experts and contain the domain-dependent specifications that are given as input to the exoskeleton specializations synthesis in order to produce exoskeletons specialized for the

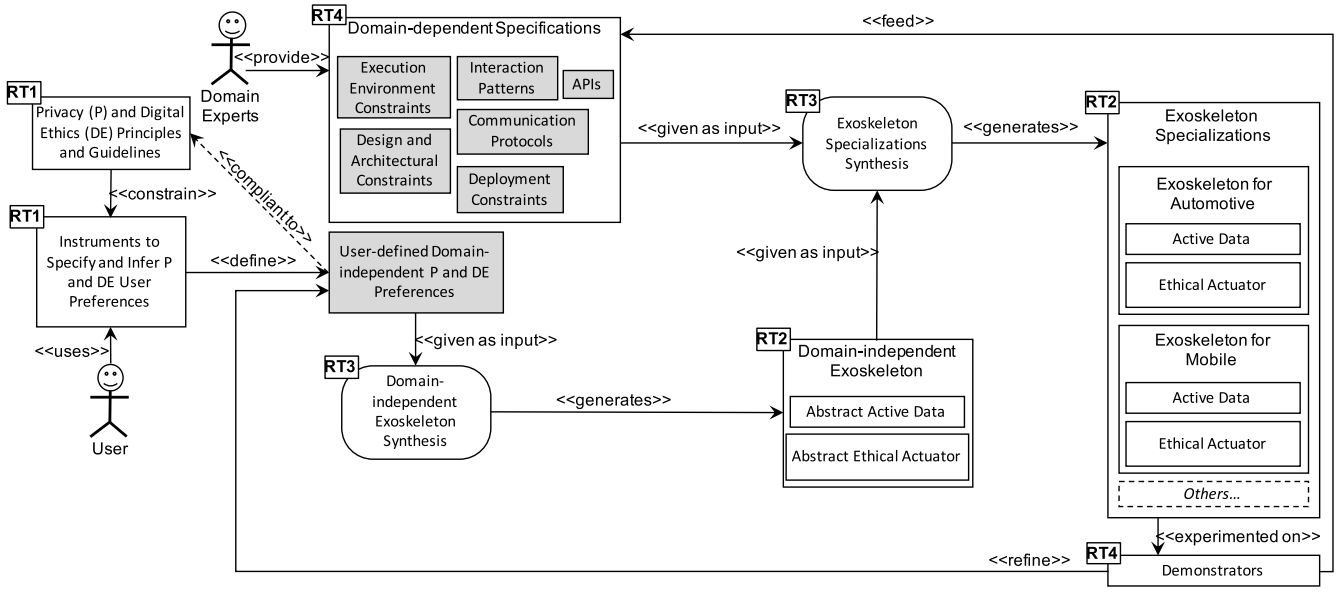


**FIGURE 1.** Overview of the **EXOSOUL** methodology.

accounted domains, e.g., automotive, mobile, etc. These specifications will be provided by following *the practical guidelines* that will be established while developing demonstrators in Research Theme 4 (RT4), as detailed in Section V-D. The demonstrators will be used to assess the research and innovation outcomes of EXOSOUL .

### *A. LOGIC THEORIES AND INNOVATIVE MECHANISMS FOR INFERRING AND SPECIFYING PRIVACY AND ETHICAL USER PREFERENCES*

To address the challenge of specifying and inferring soft ethical preferences, we will start investigating a kind of "functional morality" [56], which enables machines to autonomously assess and respond to moral challenges. The approach combines top-down and bottom-up methods. A top-down approach of encoding ethical theories (like utilitarianism or Kant's categorical imperative) in software requires making explicit ethical judgements and organize them in categories, checking which logic formalisms better represent the ethical approach devised, whether deontic – which represent duties and obligations –, Bayesian – which represents preferences –, or other. However, since people frequently departs from the norms, even those they explicitly endorse, the top-down approach should be complemented with a bottom-up approach. This approach will consider the actual behavior and the actual actions of the end-users, both in real and in simulated environments. The data collected will allow to refine and tune up the ethical principles and theories encoded by the top-down approach.

The crucial element of this research theme is the concept of functional morality. While for Wallach and Allen [56] such "functional morality" is conceived as partially autonomous behavior, we think possible to specify it in terms of dispositions. Dispositions are those properties characterized by the causal behavior of the individuals that possess them, and they are irreducibly so [27], [36]. As their name suggests, these properties dispose towards further properties: their manifestations, which occur when some conditions are met. Dispositions are intrinsic to their possessors and real, even when unmanifested. It should be clear in which sense dispositions could be considered functions: they take some value as input and manifest some behavior. But, at the same time, there is some intrinsicness in them: so, they are partially autonomous. Dispositions may also fail to manifest because of some conditions that prevent them to do so, nonetheless their possessors still have the "prevented" function. For example, a nuclear reactor contains, say, Uranium pellets that, reaching a critical mass, have the disposition to explode. If they are about to explode some sensors trigger various safety mechanisms, such as boron rods, that prevent the explosion by shutting the reactor down. These safety mechanisms prevent the manifestation of the Uranium's disposition to explode (cfr. [45], [46]) Although dispositions seem to fit perfectly within an ethical system, just a few people explored this possibility [13], [52]: moral responsibility and normativity (which both concern the bottom up and the top down method) are two key conditions for ethics and these seem to be dispositional notions [13].

In operative terms, in pursuing the top-down approach we will consider the relevant legislation of the member states (e.g., GDPR [6], ethical reference groups [8], [23]) and the normative approaches to ethics. Furthermore, we will elicit patterns for specifying privacy and ethics out of existing privacy and ethical rules defined by both the academic and industrial communities, examples of which may be found in [15], [28], [37], [40]. The ethical and privacy concerns and

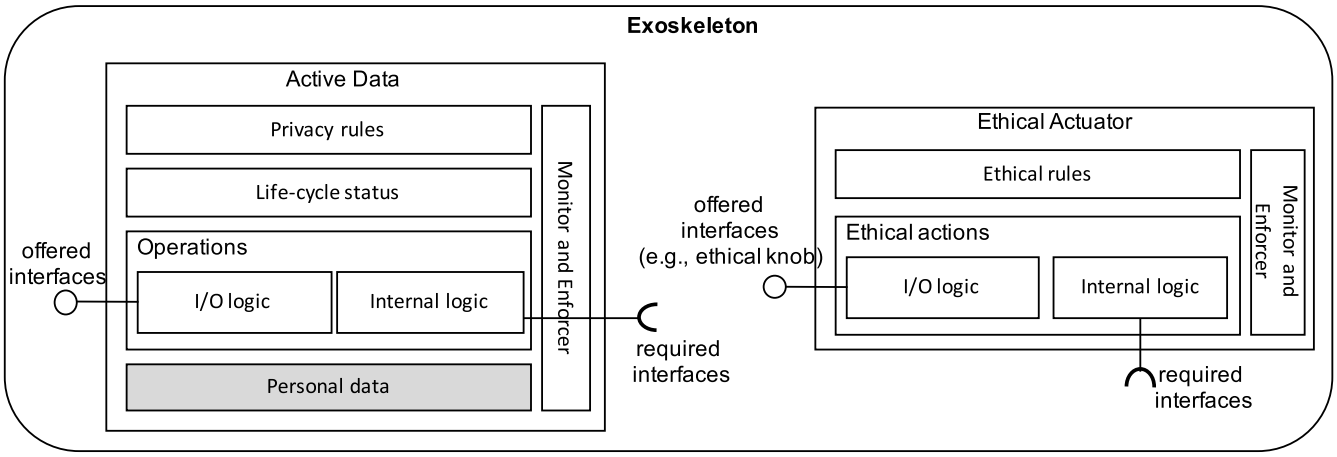


**FIGURE 2.** Exoskeleton structural overview.

their categorization, together with the elicited patterns, will be the base for refining the conceptual models and related metamodels so to enable a programmatic access to the formulation of the concepts of soft privacy and ethics and their relationships. The idea is to use instruments like questionnaires, wizards, or to build a EXOSOUL's bot and/or assistant for inferring privacy and ethical preferences. Technologies that we will investigate include Almond [24], the Google assistant App,[4] Google Duplex,[5] AI technologies like recurrent neural networks,[6] TensorFlow Extended (TFX) [20], and so on.

Regarding the bottom-up approach, we will realize it via easy-to-use feedback collection mechanisms that allow users to understand privacy rules and ethical behaviors or decisions of the system and to evaluate them. We will employ an iterative approach to the design and validation of the innovative mechanisms for inferring and specifying ethical and privacy preferences. Representative users will be in the loop at every stage. As an example, we will devise a learning "data collector" analogous to MIT's moral machine - let's call it EXOSOUL Moral Machine (EMM). It will present scenarios that comply with rules and laws but violate in various ways our more intimate and complex moral views about interacting with autonomous machines, sharing data with them, and maintaining control and independence. Reactions to such violations - in terms of graded answers rather than "yes/no" choices - will constitute data to assess the model and refine it. The questions, devised by philosophers and implemented after the EMM's feedback, will supply the learning machine with data that will contribute to define standard profiles of people's views about their own interaction with machines.

[4]https://developers.google.com/assistant/sdk/reference/library/python/
[5]https://www.wired.co.uk/article/google-io-2018-keynote-summary-recap
[6]https://en.wikipedia.org/wiki/Recurrent_neural_network

### B. EXOSKELETON DESIGN AND NEWFANGLED TECHNIQUES AND TOOLS FOR MANAGING ITS LIFE-CYCLE

Concerning the second challenge, in this section we describe the definition of the exoskeleton software architecture and of the run-time analysis mechanisms, such as monitoring and enforcement, that serve to control the exoskeleton behavior according to the specified privacy and ethics preferences.

A further subject of study regards the definition of the assumptions, i.e., required interfaces, that EXOSOUL makes on the target execution environment in order to allow the deployment and execution of its components. On the other hand, offered interfaces allow I/O interactions with the exoskeleton. The design of both required and offered interfaces will provide inputs to the definition of the practical guidelines.

Domain-independent exoskeletons and their specializations adhere to the same structure that is shown in Figure 2. An exoskeleton is composed of two parts: active data and ethical actuator.

**Active data** wrap personal data by adding the logic required to access personal data and manage their life-cycle, from creation to destruction, sharing and usage, according to the specified privacy preferences (privacy rules in Figure 2). The conformance to the privacy preferences is guaranteed by the monitor and enforcer component that makes use of the internal operations and continuously checks and updates the life-cycle status to promptly detect and correct problems before privacy-violating actions are performed. The life-cycle status keeps information like the number of data visualizations, replicas, or sharing within a (social) network, location-based information, and in general any information that allows the run-time evaluation of privacy preferences. The I/O logic, externally exposed through offered interfaces, serve to support the access (visualization, replicas, sharing) to the data by the external environment.

Differently from the canonical "passive" approach, which gives the ability to freely interact with data (e.g., from back-end systems), we opt for an active approach that, encapsulating personal data and privacy rules, monitors the interactions with the data and enforce privacy rules when needed. Exploiting the synthesis techniques defined in Research Theme 3, the active data embed a representation of the privacy preferences in terms of privacy rules that are expressed using logic formalisms, e.g., temporal logic.

For domain-independent exoskeletons, the personal data are taken from an abstract taxonomy of types of personal data, e.g., physical, physiological, genetic, mental, economic, cultural, social, etc. Similarly, the privacy rules are patterns of rules that contain abstract parameters that will be instantiated when generating an exoskeleton specialization. For instance, by referring to the automotive use case introduced in Part B1 (page 3), the personal data concern physical information about the woman and the privacy rule could be an LTL formula specifying that "a <physical data> is **shared** only **between** <start event> **and** <end event>".[7]

For exoskeleton specializations, all the data and event parameters are instantiated in variables whose value can be evaluated at run-time. For instance, the above rule is specialized into the rule: "a <pregnant status> is **shared** only **between** <entering the parking lot> **and** <exiting the parking lot>", where <pregnant status> is a Boolean variable representing the (specialized) personal data that is set by the run-time monitor the exoskeleton is equipped with. Furthermore, <entering the parking lot> and <exiting the parking lot> are events that can be caught by the monitor by exploiting the life-cycle status of the active data, e.g., GPS coordinates. To ensure a privacy rule, the monitor triggers the enforcer that uses internal operations to create, destroy, share, or use the personal data according to what is specified by the privacy rule. For instance, continuing the automotive example, the <pregnant status> of the woman is not visible anymore when her vehicle exits the parking lot.

The **ethical actuator** translates conceptual ethical principles into concrete statements that serve as the basis for ethical decision making. An ethical actuator is composed of: (i) ethical rules defined by users, (ii) a monitor, (iii) an enforcer, and (iv) ethical actions.

Analogously to privacy rules, ethical rules are parametric for domain-independent exoskeletons and afterwords instantiated for their specializations. After the synthesis phase (Research Theme 3), they are expressed through logic formalisms. For instance, continuing the automotive example, a domain-independent ethical rule could be a LTL formula specifying that "**between** <start event> **and** <end event>, **when** a <condition> holds, **then** <action> is taken".[8]

A possible specialization of it is: "**between** <entering the parking lot> **and** <exiting the parking lot>, **when** a <pregnant status> holds, **then** a <the surrounding drivers are alerted with a message>", where the <the surrounding drivers are alerted with a message> is a placeholder for the actual implementation code that sends an alert to all the surrounding drivers by exploiting the middleware-level APIs made available by the execution environment (required interfaces), e.g., the connected-vehicle infrastructure. It is worth noting that this case shows how personal privacy is strictly connected to ethics: by disclosing a personal information like this, the pregnant woman follows a utilitarian view which is related to the expectation that surrounding drivers might have a virtue soft ethic.

The monitor and the enforcer are needed to guarantee that the ethical rules will not be violated. The monitor and enforcer use also internal actions in order to accomplish their tasks. An ethical knob interface, as part of the offered interfaces, is used to either provide feedback to the users about the ethical rules application or to allow the user to take control over the execution of the ethical actions whenever some unexpected behavior is raising. Furthermore, exoskeletons will be designed so to achieve robustness and security by exploiting state of the art solutions.

### C. EXOSKELETON SYNTHESIS

In this section we describe the definition and the realization of automated synthesis methods for the generation of: (i) a domain-independent exoskeleton starting from the user's ethical and privacy preferences, and (ii) a domain-specific specialization of the domain-independent exoskeleton from the inputs provided by domain experts. These inputs regard information that are required to produce the code of the specialized exoskeleton, and package it as required by the target execution environment.

By referring to the exoskeleton structural overview shown in Figure 2, the user's ethical and privacy preferences will be transformed into ethical and privacy rules, respectively. This transformation step concerns also the synthesis of the status variables that will allow the synthesized exoskeleton to control the active data life-cycle via monitoring and enforcement. Furthermore, starting from the ethical and privacy rules, code templates are generated for: (i) the enforcer and its related monitor, (ii) the active data operations, and (iii) the ethical actions, hence synthesizing the domain-independent exoskeleton. Domain specialization if performed by using inputs from the domain experts that account for: (i) design and architectural constraints, communication protocols, which drive the way the specialized exoskeleton will communicate and interact with the other components in the target system;

[7]The actual LTL formula is: $([\,](Q\ \&\ !R \rightarrow (!R\ W(P\ \&\ !R))))\&(<>R \rightarrow (!PUR))\&([\,](Q \rightarrow [\,](!P)))$, where $P$ is "<physical data> is **shared**", $Q$ is "<start event>", and $R$ is "<end event>". The formula is obtained via conjunction of Existence (between $Q$ and $R$), Absence before $R$, and Absence after $Q$ specification patterns [28].

[8]The actual LTL formula is: $[\,]((Q \wedge \neg R \wedge \diamond R) \rightarrow (P \rightarrow (\neg R \cup (S \wedge \neg R))) \cup R)$, where $P$ is <condition>, $S$ is <action>, $R$ is <start event>, and $Q$ is <end event>. The formula is obtained via the specification pattern response (between $Q$ and $R$) [28].

(ii) execution environment and deployment constraints that allow the synthesis to derive how the specialized exoskeleton has to be packaged, deployed, and enacted; (iii) system component APIs used to match the required exoskeleton interfaces.

The approach just described is extremely challenging since it has to cope with the complexity of representing and enforcing ethical and privacy rules. However, we can base on our expertise on the prevention of interaction mismatches. Indeed, in our previous work [16]–[19], [39], [55] we exploited architectural specifications, including interaction and communication patterns, APIs, etc., to automatically generate integration and coordination code for the components forming a target distributed system. The code is generated out of an intermediate behavioral model of a connector that is automatically synthesized from the system architectural specification. Then, by a further synthesis step, the connector model is specialized to behave as a centralized coordinator [55] or a distributed one [19], a mediator [16], [39], or a distributed enforcer [17], [18], depending on the coordination issues, the interaction protocol mismatches to be solved, and the nature of the target system. As a consequence of the adopted approaches, all these integration and coordination artifacts embed monitoring logic.

By exploiting the outcomes of Research Theme 1, our research work will consider specifications of user's ethical and privacy preferences expressed in terms of a structured natural language, e.g., based on Structured English Grammar [15] (SEG). We will define a collection of mappings from the SEG to the identified logic formalisms in order to allow the synthesis to transform the ethical and privacy preferences into privacy and ethical rules. To this aim, we will also build a catalogue of specification patterns by coordinating existing pattern catalogues [21], [28], [37], [40]. The existing catalogues are general purpose catalogues defined with the aim to support engineers in the verification of systems. We will start from these catalogues and we will define catalogues that are tailored to ethics and privacy by using suitable alignment procedures via literature review, gap analysis, and pattern elicitation.

Concerning the synthesis of the domain-independent exoskeleton, we will follow a model-to-model generative approach. For instance, ground terms of a privacy rule may constitute the status variables encoding the active data lifecycle. The model that expresses the semantics of the generated rules (e.g., automata-based models, sequence charts, Markov chains) is analyzed to produce a behavioral model of the monitor. The model specifies the (sequence of) events to monitor at run-time, and the related information needed to assure the execution flows specified by the rule. It can specify the events that lead to interactions that violate the rule and controllability/uncontrollability of these events. The model of the monitor is then analyzed to synthesize a model of the enforcer that will encode only the execution flows that assure the rule. It can perform backwards error propagation on the events that lead to violating interactions, depending on their controllability/uncontrollability. Concerning I/O logic, code templates will be generated to handle the deployment and enactment of the active data part. Starting from the ethical rules, we will follow a similar approach for the synthesis of the ethical actuator part.

For the synthesis of the domain-specific exoskeleton, we will follow a model-to-code generative approach. As anticipated in Section IV, a crucial aspect here concerns the huge amount and the heterogeneity of the information about the target execution environment that are needed to produce the actual (skeleton) code of the specialized exoskeleton components. Assuming that domain experts will provide this information complete and in a form that is suitable for EXOSOUL is unrealistic. Thus we will also study inference mechanisms to discover characteristics of the execution environment, which complement the analysis of the, e.g., automotive and mobile domains and the experimentation with manually-coded exoskeletons within Research Theme 4. In our previous work [22], we defined a method to automatically infer a component's interaction protocol out of a specification of its interface. In EXOSOUL, we have to account for the automated elicitation of different domain-dependent specifications, not only interaction protocols.

### D. DEMONSTRATORS AND PRACTICAL GUIDELINES

This research theme has the twofold objective of continuously experimenting the research outcome to validate and guide the performed research, and to steam out of it practical guidelines for companies and organizations willing to adopt EXOSOUL .

We plan to exercise in the automotive and mobile domains. This will ensure that EXOSOUL builds on real characteristics of these highly-evolving domains and can deliver practical results to implement proof-of-concept demonstrators. The two scenarios will serve as testbeds and benchmarks for the solutions developed, resulting in rapid feedback for steering the research activities in EXOSOUL. On the one hand, experiments will validate exoskeletons against specified privacy and ethical preferences. On the other hand, as shown in Figure 1 (see the <<refine>> and <<feed>> arrows from the Demonstrators block), the experiments will be used to understand users behavior, which will enable the refinement of the user profile to possibly adjust her ethical and privacy preferences based on her actions in the field. Experiments will also influence the development of recommendations in order to provide interested stakeholders (such as platform vendors, IT big players, software producers, domain experts) with practical guidelines to enable the adoption of exoskeletons. For example, the practical guidelines will concern the definition of architectural, protocol, and development constraints that need to be accepted and satisfied by platform vendors that wish to employ EXOSOUL, provided that some (business) opportunities have been identified. For example, OEMs shall accept that the exoskeleton software of the driver and/or passenger(s) can be deployed on board and can interact with the vehicle, or that a new data format needs to be accounted for by some mobile apps. This may require to define a set

of application- and middleware-level APIs that should be exposed by the execution environment to let synthesized exoskeletons actually achieve their enforcing tasks.

Concerning the automotive domain, we will exploit our connection in FCA (Fiat Chrysler Automobiles) towards the AdaptIVe FP7 EU project[9] (FP7-ICT-2013.6.5, "Automated Driving Applications and Technologies for Intelligent Vehicles"). We plan to use a driving simulator whose level of automation is "Conditional Automation", i.e., level 3 according to the classification developed through the Society of Automotive Engineers (SAE) International, for defining driving automation for motor vehicles.[10] Conditional automation means that we will be able to simulate scenarios of automated driving in dense freeway traffic (low speeds) and in limited areas/roads, where the driver can possibly take over after warning. For instance, the automated parking scenario introduced in Section III conforms to this kind of scenarios.

For what concerns the mobile domain, we will mainly experiment in the Android ecosystem. Android is expected to continue to capture roughly 85% of the worldwide smartphone volume (Android and Apple cover 99.9%).[11] With a recent tweet by the official Google account, Google announces that Android now has 2 Billion Monthly active users.[12]

## VI. CONCLUSION

Our long-term vision is twofold. On the one hand, we want to promote a digital world where digital actors and humans are in better balance of forces. Whenever needed, citizens will be able to exploit software exoskeletons that, reflecting their own ethics and privacy, empower and protect them by mediating their interactions with the digital world. They will be responsible for the consequences of their choices, even if mediated by autonomous technologies (a sensible issue at this stage of autonomous vehicles R&D), and they will be able to enforce their own ethical and privacy values in the behavior of the smart and autonomous systems they daily use and interact with. Users will not anymore need to passively accept the opaque behavior of the software technologies they use, rather EXOSOUL will provide them with effective means to mold these technologies into the shape that better reflects their human perception and needs. This will permit to achieve "the principle of human dignity, understood as the recognition of the inherent human state of being worthy of respect, must not be violated by autonomous technologies".[13] At the same time, we want to influence the stakeholders involved in the development and operation of networked applications, as well as the stakeholders involved in the creation of standards and laws relative to digital ethics. In the long term, EXOSOUL has the ambition to revolutionize the balance of power in the digital world, introducing more symmetry among the various actors, notably software companies and citizens.

Furthermore, EXOSOUL will create new opportunities for existing and new companies in the field of societal-friendly applications. Europe is much sensible to the theme of privacy, data protection, and moral and ethical aspects in the digital world, and through EXOSOUL it can really become the scientific and technological leader of the future ethically-aware systems. In fact, EXOSOUL will deliver the first concrete contribution to an ethical approach *to regulate the digital world* in line with the goals of the European Data Protection Supervisor strategy 2015-2019 [5].

[9] http://www.adaptive-ip.eu
[10] https://www.sae.org
[11] https://www.gartner.com/newsroom/id/3859963
[12] https://twitter.com/Google/status/864890655906070529
[13] https://ec.europa.eu/research/ege/pdf/ege_ai_statement_2018.pdf

## ACKNOWLEDGMENTS

The authors would like to thank the entire multi-disciplinary team of the EXOSOUL@univaq project (http://exosoul.disim.univaq.it) for enlightening debates and joint work on digital ethics for autonomous systems.

## REFERENCES

[1] *Gartner Says Worldwide Sales of Smartphones Recorded First Ever Decline During the Fourth Quarter of 2017*. Accessed: Feb. 19, 2019. [Online]. Available: https://www.gartner.com/newsroom/id/3859963
[2] *The IRMA Project*. Accessed: Feb. 19, 2019. [Online]. Available: https://privacybydesign.foundation/irma-en/
[3] *Virtue Ethics*. Accessed: Feb. 19, 2019. [Online]. Available: https://plato.stanford.edu/entries/ethics-virtue/
[4] (2018). *The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems*. [Online]. Available: https://standards.ieee.org/develop/indconn/ec/autonomous_systems.html
[5] *2017 Annual Report—Data Protection and Privacy in 2018: Going Beyond the GDPR*. (2018). [Online]. Available: https://edps.europa.eu/sites/edp/files/publication/18-03-15_annual_report_2017_en.pdf
[6] *General Data Protection Regulation*, Eur. Commission, Brussels, Belgium, 2018.
[7] (2018). *The European Commission's High-Level Expert Group on Artificial Intelligence, Draft Ethics Guidelines for Trustworthy Ai*. [Online]. Available: https://ec.europa.eu/digital-single-market/en/news/nhave-your-say-european-expert-group-seeks-feedback-draft-nethics-guidelines-trustworthy
[8] (2018). *European Group on Ethics in Science and New Technologies*. Statement on Artificial Intelligence, Robotics and 'Autonomous' Systems. [Online]. Available: https://ec.europa.eu/research/ege/pdf/egeaistatement2018.pdf
[9] (2018). *Partnership on AI*. [Online]. Available: https://www.partnershiponai.org/
[10] (2018). *Why The Netherlands is the Globe's top Location for Selfdriving Cars*. https://www.consultancy.eu/news/1098/why-the-netherlands-is-the-globes-top-location-for-self-driving-cars
[11] *Gartner Top 10 Strategic Technology Trends for 2019*. (2019). [Online]. Available: https://gtnr.it/2CJJYGp
[12] (2019). *The State of Autonomous Legislation in Europe*. [Online]. Available: https://autovistagroup.com/news-and-insights/state-autonomous-legislation-europe
[13] R. Anjum, S. Noer Lie, and S. Mumford, "Dispositions and ethics," in *Forthcoming in Anjum and Mumford (ed), What Tends to Be: The Philosophy of Dispositional Modality*. Abingdon, U.K.: Routledge, 2018.
[14] (2018). *Association for Computing Machinery U.S. Public Policy Council (USACM)*. Statement on Algorithmic Transparency and Accountability. [Online]. Available: https://www.acm.org/binaries/content/assets/public-policy/2017usacmstatementalgorithms.pdf
[15] M. Autili, L. Grunske, M. Lumpe, P. Pelliccione, and A. Tang, "Aligning qualitative, real-time, and probabilistic property specification patterns using a structured English grammar," *IEEE Trans. Softw. Eng.*, vol. 41, no. 7, pp. 620–638, Jul. 2015.
[16] M. Autili, P. Inverardi, F. Mignosi, R. Spalazzese, and M. Tivoli, "Automated synthesis of application-layer connectors from automata-based specifications," in *Proc. 9th Int. Conf. Lang. Automata Theory Appl. (LATA)*, 2015, pp. 3–24.

[17] M. Autili, P. Inverardi, and M. Tivoli, "Automated synthesis of service choreographies," *IEEE Softw.*, vol. 32, no. 1, pp. 50–57, Jan. 2015.
[18] M. Autili, P. Inverardi, and M. Tivoli, "Choreography realizability enforcement through the automatic synthesis of distributed coordination delegates," *Sci. Comput. Program.*, vol. 160, pp. 3–29, Aug. 2018.
[19] M. Autili, L. Mostarda, A. Navarra, and M. Tivoli, "Synthesis of decentralized and concurrent adaptors for correctly assembling distributed component-based systems," *J. Syst. Softw.*, vol. 81, no. 12, pp. 2210–2236, 2008.
[20] D. A. Baylor *et al.*, "TFX: A tensorflow-based production-scale machine learning platform," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*. New York, NY, USA: ACM, 2017, pp. 1387–1395.
[21] P. Bellini, P. Nesi, and D. Rogai, "Expressing and organizing real-time specification patterns via temporal logics," *J. Syst. Softw.*, vol. 82, no. 2, p. 183–196, 2009.
[22] A. Bertolino, P. Inverardi, P. Pelliccione, and M. Tivoli, "Automatic synthesis of behavior protocols for composable web-services," in *Proc. ESEC/FSE*, 2009, pp. 141–150.
[23] J. P. Burgess, L. Floridi, A. Pols, and J. van den Hoven. (2018). *Towards A Digital Ethics–EDPS Ethics Advisory Group*. [Online]. Available: https://edps.europa.eu/sites/edp/files/publication/18-01-25_eag_report_en.pdf
[24] G. Campagna, R. Ramesh, S. Xu, M. Fischer, and M. S. Lam, "Almond: The architecture of an open, crowdsourced, privacy-preserving, programmable virtual assistant," in *Proc. 26th WWW*, 2017, pp. 341–350.
[25] G. Contissa, F. Lagioia, and G. Sartor, "The ethical knob: Ethically-customisable automated vehicles and the law," *Artif. Intell. Law*, vol. 25, no. 3, pp. 365–378, 2017.
[26] V. Cortellessa, A. D. Marco, and P. Inverardi. *Model-Based Software Performance Analysis*, 1st ed. Springer, 2011.
[27] D. Donati. (2018). *No Time for Powers*. [Online]. Available: http://eprints.nottingham.ac.uk/id/eprint/51347
[28] M. B. Dwyer, G. S. Avrunin, and J. C. Corbett, "Property specification patterns for finite-state verification," in *Proc. ICSE*. New York, NY, USA: ACM, 1999, pp. 411–420.
[29] EDPS. (2015). *Leading by Example, The EDPS Strategy 2015-2019*. [Online]. Available: https://edps.europa.eu/sites/edp/files/publication/15-07-30_strategy_2015_2019_update_en.pdf
[30] EDPS. (2015). *Opinion 4/2015, Towards a New Digital Ethics–Data, Dignity and Technology*. [Online]. Available: https://edps.europa.eu/sites/edp/files/publication/15-09-11_data_ethics_en.pdf
[31] J. Larus *et al.* (2018). *When Computers Decide: European Recommendations on Machine-Learned Automated Decision Making*. [Online]. Available: http://www.acm.org/binaries/content/assets/public-policy/ie-euacm-adm-report-2018.pdf
[32] T. Evas, C. Rohr, F. Dunkerley, and D. Howarth, *A Common EU Approach to Liability Rules and Insurance for Connected and Autonomous Vehicles*. Luxembourg City, Luxembourg: Publications Office of the European Union, 2018. doi: 10.2861/282501.
[33] L. Floridi, "Soft ethics and the governance of the digital," *Philosophy Technol.*, vol. 31, no. 1, pp. 1–8, Mar. 2018.
[34] S. García, C. Menghi, P. Pelliccione, T. Berger, and R. Wohlrab, "An architecture for decentralized, collaborative, and autonomous robots," in *Proc. IEEE Int. Conf. Softw. Archit. (ICSA)*, Apr. 2018, pp. 75–84.
[35] J. Gogoll and J. F. Müller, "Autonomous cars: In favor of a mandatory ethics setting," *Sci. Eng. Ethics*, vol. 23, no. 3, pp. 681–700, Jun. 2017.
[36] S. Gozzano, "L'essenzialismo scientifico e il mentale," *Rivista di Filosofia*, vol. 103, no. 2, pp. 201–226, 2012.
[37] L. Grunske, "Specification patterns for probabilistic quality properties," in *Proc. ICSE*, W. Schäfer, M. B. Dwyer, and V. Gruhn, Ed. New York, NY, USA: ACM, 2008, pp. 31–40.
[38] P. Inverardi, "The european perspective on responsible computing," *Commun. ACM*, vol. 62, no. 4, p. 64, Mar. 2019.
[39] P. Inverardi and M. Tivoli, "Automatic synthesis of modular connectors via composition of protocol mediation patterns," in *Proc. 35th Int. Conf. Softw. Eng. (ICSE)*, 2013, pp. 3–12.
[40] S. Konrad and B. H. C. Cheng, "Real-time specification patterns," in *Proc. ICSE*. New York, NY, USA: ACM, 2005, pp. 372–381.
[41] B. Krishnamurthy and C. E. Wills, "Characterizing privacy in online social networks," in *Proc. 1st Workshop Online Social Netw. (WOSN)*. New York, NY, USA: ACM. 2008, pp. 37–42.
[42] I. Malavolta, P. Lago, H. Muccini, P. Pelliccione, and A. Tang, "What industry needs from architectural languages: A survey," *IEEE Trans. Softw. Eng.*, vol. 39, no. 6, pp. 869–891, Jun. 2013.
[43] I. Malavolta, H. Muccini, P. Pelliccione, and D. Tamburri, "Providing architectural languages and tools interoperability through model transformation technologies," *IEEE Trans. Softw. Eng.*, vol. 36, no. 1, pp. 119–140, Jan. 2010.
[44] (Aug. 2018). *Moral Machine MIT*. http://moralmachine.mit.edu/
[45] S. Mumford, *Dispositions*. Oxford, U.K.: Oxford Univ. Press, 1998.
[46] S. Mumford and R. L. Anjum. *Getting Causes From Powers*. New York, NY, USA: Oxford Univ. Press, 2011, p. 272.
[47] P. Pelliccione, P. Inverardi, and H. Muccini, "CHARMY: A Framework for Designing and Verifying Architectural Specifications," *IEEE Trans. Softw. Eng.*, vol. 35, no. 3, pp. 325–346, May 2009.
[48] P. Pelliccione *et al.*, "Automotive architecture framework: The experience of volvo cars," *J. Syst. Archit.*, vol. 77, pp. 83–100, Jun. 2017.
[49] P. Samarati and S. C. de Vimercati, "Access control: Policies, models, and mechanisms," in *Foundations of Security Analysis and Design*, R. Focardi and R. Gorrieri, Ed. Berlin, Germany: Springer, 2001, pp. 137–196.
[50] G. L. Scoccia, I. Malavolta, M. Autili, A. Di Salle, and P. Inverardi, "User-centric android flexible permissions," in *Proc. IEEE/ACM 39th Int. Conf. Softw. Eng. Companion (ICSE-C)*, May 2017, pp. 365–367.
[51] G. L. Scoccia, S. Ruberto, I. Malavolta, M. Autili, and P. Inverardi, "An investigation into Android run-time permissions from the end users' perspective," in *Proc. 5th IEEE/ACM Int. Conf. Mobile Softw. Eng. Syst. (MOBILESoft)*, May 2018, pp. 45–55.
[52] M. Smith, D. Lewis, and M. Johnston, "Dispositional theories of value," *Proc. Aristotelian Soc.*, vol. 63, no. 1, pp. 89–174, 1989.
[53] L. R. Sütfeld, R. Gast, P. König, and G. Pipa, "Using virtual reality to assess ethical decisions in road traffic scenarios: Applicability of value-of-life-based models and influences of time pressure," *Frontiers Behav. Neurosci.*, vol. 11, no. 122, pp. 1–13, 2017.
[54] S. M. Thornton, S. Pan, S. M. Erlien, and J. C. Gerdes, "Incorporating ethical considerations into automated vehicle control," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 6, pp. 1429–1439, Jun. 2017.
[55] M. Tivoli and P. Inverardi, "Failure-free coordinators synthesis for component-based architectures," *Sci. Comput. Program.*, vol. 71, no. 3, pp. 181–212, 2008.
[56] W. Wallach and C. Allen, *Moral Machines: Teaching Robots Right from Wrong*. New York, NY, USA: Oxford Univ. Press, 2010.
[57] P. Wijesekera *et al.*, "Contextualizing privacy decisions for better prediction (and protection)," in *Proc. CHI Conf. Hum. Factors Comput. Syst. (CHI)*. New York, NY, USA: ACM, 2018, p. 268.
[58] C. Yang *et al.*, "An industrial case study on an architectural assumption documentation framework," *J. Syst. Softw.*, vol. 134, p. 190–210, Dec. 2017.

**MARCO AUTILI** is currently an Assistant Professor with the DISIM, University of L'Aquila. He is (has been) involved in many EU and Italian projects contributing to research and development activities as well as management and coordination activities. His main research interests are in software engineering and include the (from theory to practice) application of methods to the verification, analysis and automatic synthesis of distributed systems, software architectures with particular reference to service-oriented architectures, and application of context-oriented programming techniques to the development of adaptable (mobile) applications. He published several papers in journals, international conferences, and workshops in these topics. He is on the Program Committee of several international conferences and workshops and also a Reviewer of journals, including *Science of Computer Programming*, *Software and Systems Modeling*, the *Journal of Systems and Software*, *Automated Software Engineering*, and a number of transactions. More information is available at http://people.disim.univaq.it/marco.autili.

**DAVIDE DI RUSCIO** is currently an Associate Professor with the DISIM, University of L'Aquila. His main research interests include several aspects of software engineering, open-source software, and model-driven engineering (MDE), including domain-specific modeling languages, model transformation, model differencing, and model evolution. He has published more than 130 papers in various journals, conferences, and workshops on these topics. He is a member of the Steering Committee of the International Conference on Model Transformation (ICMT), the Software Language Engineering (SLE) Conference, the Seminar Series on Advanced Techniques & Tools for Software Evolution (SATTOSE), the Workshop on Modelling in Software Engineering at ICSE (MiSE), and the International Workshop on Robotics Software Engineering (RoSE). He is on the Editorial Board of the *International Journal on Software and Systems Modeling* (SoSyM), *The Journal of Object Technology*, and the *IET Software* journal. Since 2006, he has been working on different European and Italian research projects by contributing the application of MDE concepts and tools in several application domains (e.g., service-based software systems, autonomous systems, open-source software systems, and hybrid polystore systems). More information is available at http://people.disim.univaq.it/diruscio.

**PAOLA INVERARDI** is currently a Full Professor with the DISIM, University of L'Aquila, where she is also a Rector. Her research interests include the field of the application of rigorous techniques to the development of software systems, including software specification and verification of concurrent and distributed systems, software architectures, and synthesis of connectors. Her current research interests include the field of software architectures addressing the verification and analysis of software architecture properties, both behavioral and quantitative, the synthesis of correct by construction distributed systems, and responsible computing. She is a member of the EUACM and the Academia Europaea. She received the Honorary Doctorate from Mälardalen University and Shibaura University. More information is available at http://people.disim.univaq.it/~inverard.

**PATRIZIO PELLICCIONE** received the Ph.D. degree from the University of L'Aquila, in 2005. He is currently an Associate Professor with the Department of Computer Science and Engineering, Chalmers University of Technology, the University of Gothenburg, Sweden, and the DISIM, University of L'Aquila, Italy. He is very active in European and national projects. His research interests include software engineering, software architectures modeling and verification, autonomous systems, and formal methods. He has coauthored more than 130 publications in journals and international conferences and workshops in these topics. In his research activity, he has collaborated with several industries, such as Volvo Cars, Volvo AB, Ericsson, Jeppesen, Axis Communication, Thales Italia, Selex Marconi Telecommunications, Siemens, Saab, and TERMA. In 2014, he received the Docent in Software Engineering from the University of Gothenburg. He has been on the program committees for several conferences. He is also a Reviewer for top journals in the software engineering domain. More information is available at http://www.patriziopelliccione.com.

**MASSIMO TIVOLI** is currently an Associate Professor with the DISIM, University of L'Aquila. His main research interests include the definition and application of formal software engineering methods to the development of distributed software systems. In particular, one of his research topics include the automated synthesis of correct-by-construction connectors for the correct assembly of distributed component-based systems, focusing on the production of both centralized and distributed connectors. Other research topics include the development of dependable and adaptable systems, automated methods to learn the interaction protocol performed by a service directly out of its signature description, and automated choreography synthesis approaches. He has been involved in several European and national projects, among which the European H2020 CHOReVOLUTION Project for which he was a Scientific Coordinator. He is also a Reviewer for several leading international peer-reviewed journals. He participated/participates on program committees of several relevant conferences on software architectures and software engineering. He has been the Program Chair of the international symposium, CBSE2013. More information is available at http://people.disim.univaq.it/massimo.tivoli.