

Acceptability of Symbiotic Artificial Intelligence: Highlights from the FAIR project

Francesca Alessandra Lisi^{1,*}, Antonio Carnevale², Abeer Dyoub¹, Antonio Lombardi², Piero Marra³ and Lorenzo Pulito⁴

¹University of Bari Aldo Moro, DiB Dept., via E. Orabona 4, Bari, 70125, Italy

²University of Bari Aldo Moro, DIRIUM Dept., Piazza Umberto I, Bari, 70121, Italy

³University of Bari Aldo Moro, LAW Dept., Piazza C. Battisti 1, Bari, 70121, Italy

⁴University of Bari Aldo Moro, DJSGE Dept., Via Duomo 259, Taranto, 74123, Italy

Abstract

In this work we report the highlights of the work done at the University of Bari within the FAIR project and concerning the acceptability of Symbiotic Artificial Intelligence.

Keywords

Symbiotic AI, AI Ethics, Trustworthy AI, Philosophical foundations of AI

1. Introduction

The notion of *symbiosis* originated in the 19th century to indicate a relationship between two taxonomically separate life forms that nevertheless give rise to a single organism. Life forms in a symbiotic relationship are not isolated but coexist in ways that are more or less essential to their survival and development. The first to advocate a symbiosis between humans and machines was J.C.R Licklider in 1960 [1]. In his view, this kind of symbiosis would allow the computer to become an active part of the thinking process that leads to resolving technical problems and not just an executor of solutions thought up beforehand. Licklider was mainly thinking of human-computer interfaces that would allow greater real-time collaboration and shorten the distance between human and machine language. He was pointing to a road that has since been successfully travelled, bringing us to the so-called *Symbiotic Artificial Intelligence* (SAI). Human-AI symbiosis promises to boost human-machine collaboration and socio-technical teaming, with mutually beneficial relationships, by augmenting (and valuing) human cognitive abilities rather than replacing them [2]. In particular, *socio-technical teaming* refers to the collabora-

tive partnership between humans and machines within a broader social and technological context, where the focus is not on a substantial peer-to-peer relationship but on integrating technology into human-centric processes and systems. In this context, symbiosis involves humans and machines working together as a cohesive unit, each playing a specific role and contributing to the team's overall performance. On one hand, humans provide the cognitive and emotional capabilities necessary for creativity, empathy, ethical decision-making, and adaptability. On the other hand, machines offer computational power, data processing, and automation capabilities that can handle repetitive and data-intensive tasks efficiently.

When applied to AI, the concept of symbiosis becomes more complex, posing a whole series of foundational questions. Addressing these questions is one of the goals of the research done by the University of Bari (together with INFN) within the project Future AI Research (FAIR). In particular, the *acceptability* of SAI is the subject of research for our investigation within a dedicated work package (WP 6.5) of FAIR. Acceptability involves value alignment between AI and humans. It is related, e.g., to understanding AI decisions, the algorithmic bias, the respect of privacy policies for data collected by AI systems, the struggle between security ensured by AI systems and fundamental freedoms, the mitigation of possible safety and health risks. In FAIR, studies on the acceptability of SAI adopt an interdisciplinary approach involving researchers in AI, Law, and Philosophy.

In this paper, we briefly report the main achievements of our research on ethical and legal acceptability of SAI in the 1st year of the project (Sections 2-3) and outline the steps needed to go from general principles to operational definitions for ethical acceptability (Section 4). Section 5 concludes the paper with final remarks.

Ital-IA 2024: 4th National Conference on Artificial Intelligence, organized by CINI, May 29-30, 2024, Naples, Italy

*Corresponding author.

✉ francesca.lisi@uniba.it (F. A. Lisi); antonio.carnevale@uniba.it (A. Carnevale); abeer.dyoub@uniba.it (A. Dyoub); antonio.lombardi@uniba.it (A. Lombardi); piero.marra@uniba.it (P. Marra); lorenzo.pulito@uniba.it (L. Pulito)

🆔 0000-0001-5414-5844 (F. A. Lisi); 0000-0003-2538-5579

(A. Carnevale); 0000-0003-0329-2419 (A. Dyoub);

0000-0003-1803-5423 (A. Lombardi); 0009-0003-6365-2129

(P. Marra); 0009-0000-3979-8716 (L. Pulito)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

2. Ethical acceptability of SAI

The philosophical approach to AI is contributing to the debate on the identification and analysis of the ethical implications of algorithms. We have continued the investigation aiming to build the proposal of a methodological framework grounded in process-oriented evaluations to assess the human-centricity and acceptability of SAIs together with their societal benefit.

The research carried out concerned two different scientific lines:

Questioning the notion of “symbiosis” in SAI systems. The research focused mainly on the meaning of “symbiosis” and its applicability to AI [3]. To this end, preliminary research has been carried out on the transformation of the concept of intelligence in the history of ideas [4]. In several internal meetings, the notion of symbiosis was explored both from a biological and phenomenological point of view, with reference to the key recent AI-driven technological developments (AI and drones, AI and robotics, LLM, ML, etc.).

Assessing the ethical impact of SAI in terms of acceptability and human-centricity. Defining the fundamental conceptual stages of a methodology for evaluating AI systems involves comparing and studying a series of international regulatory frameworks – *inter alia* AI HLEG, Ethics Guidelines for Trustworthy AI (2018-19). We have outlined a model with different fundamental steps: (a) onto-epistemic foundation of the method; (b) screening; (c) risk evaluation; (d) impact assessment. Now, we need to work within each step to refine procedures and metrics further.

The efforts in this direction have led to a joint paper presented at the BEWARE workshop organized in Rome within the 22nd International Conference of the Italian Association for Artificial Intelligence (AI*IA 2023) [5], an article accepted for publication in the journal *Intelligenza Artificiale* [6], and different book chapters in the final stages of publication [7], [8].

3. Legal acceptability of SAI

In line with the ethical and philosophical considerations on symbiosis, moving from the perspective of human-machine interaction to a procedural model of construction and assessment of SAI decision, within a legal methodology theory we have identified the first legal pragmatic conditions of algorithmic decision-making, such as that of the *significant human control*, a notion borrowed from the international debate within the UN on autonomous weapons. In this way, symbiosis translates also a techno-procedural legal principle capable of formalizing a human-centric value where persons do not remain behind technological development and society

but are an integral part of the same evolutionary process and are responsible for it. We think that this approach is keeping with the provisions, *ex multis*, of memorandum no. 38 of the Proposal for a EU Reg. on artificial intelligence. A procedural condition ensures the fairness and transparency of decision-making and it allows recipients to understand and respect the decision itself. Indeed, in law, it is not sufficient the content of the decision, but also its enforcement. Thus, the effectiveness remains a constitutive element of legality [9].

Furthermore, some legal issues raised by the interaction between humans and AI were addressed in some areas of law (such as those that most require judgments of predictive type, like the assessment of dangerousness aimed, for example, at commensurate punishment and/or granting alternative measures). It has been so possible to observe and identify some essential conditions that should be taken into account in designing AI systems in this field, necessary to promote the symbiosis between humans and AI as well as to improve the trustworthiness, fairness and efficiency of the interaction (for example, enriching the methods of responding to the crime in compliance with the fundamental principles of proportionality and dignity of the person, realizing the requests for individualization of the punishment) [10].

Finally, we would like to mention that, the European legal framework for AI gives minimal consideration to regulating AI based technologies where there is a reciprocal relationship between human and machine (symbiosis). The research field of symbiotic AI is technologically challenging. In [11], we have undertaken a foundational study with the aim of conceptualizing and designing a comprehensive symbiotic approach to AI, with the goal of producing fair, legitimate, and effective outcomes while ensuring their ethical and legal acceptability. This theoretical research is expected to influence the development of Symbiotic AI systems and technological governance through model assessment.

4. Towards Operational Definitions of Ethical Acceptability of SAI

The ethical implications of Human-AI symbiosis are multifaceted and complex. Thus, it has become increasingly paramount to take in consideration the ethical issues surrounding SAI development, deployment, and impact. The concept of ‘SAI Ethics’ offers a nuanced perspective that emphasizes the harmonious coexistence and collaboration between humans and AI systems. Operationalizing SAI Ethics involves translating abstract ethical principles and values into concrete guidelines and practices that govern every stage of the AI lifecycle, including

data collection, algorithm design, model training, evaluation, and deployment [12]. It requires a multidisciplinary approach, involving collaboration between computer scientists, ethicists, policymakers, and other stakeholders to ensure their alignment with societal values and human well-being, and to foster harmony and mutual benefit between humans and machines.

4.1. Operationalizing SAI Ethics

From a practical perspective, operationalizing SAI Ethics requires the establishment of governance frameworks, standards, and regulations to govern the responsible development, deployment, and use of AI technologies. This includes the development of ethical guidelines, codes of conduct, and best practices to guide AI practitioners and organizations in navigating ethical dilemmas and decision-making processes [13]. These tools should be domain specific. Moreover, fostering interdisciplinary collaboration and stakeholder engagement is essential to ensure that ethical considerations are adequately addressed and that AI technologies serve the broader societal interest.

One key aspect of operationalizing SAI Ethics is the development of robust frameworks and methodologies for ethical risk assessment and mitigation. This involves identifying potential ethical risks associated with AI systems, such as bias, discrimination, privacy violations, and unintended consequences, and implementing strategies to address these risks proactively [14]. Thus, it is important to design algorithms and systems that are transparent, interpretable, and accountable, enabling stakeholders to understand how AI decisions are made and to detect and rectify ethical issues when they arise. Here we would like to highlight the role of logic programming for designing such models [15]. Additionally, operationalizing SAI Ethics requires ongoing monitoring and evaluation of AI systems in real-world contexts to ensure that they continue to operate ethically and responsibly throughout their lifecycle. From a technical perspective, operationalization should focus on *human-centricity* through the development of AI systems that are transparent, interpretable, and accountable. This entails implementing mechanisms for explainability and interpretability, allowing users to understand how AI algorithms make decisions and providing insights into their underlying processes. Techniques such as model interpretability, transparency tools, and algorithmic audits enable stakeholders to scrutinize AI systems and identify potential biases, errors, or unintended consequences. Additionally, ensuring the robustness and reliability of AI systems through rigorous testing, validation, and verification processes is essential to minimize the risk of harmful outcomes and instil confidence in their use.

Furthermore, operationalizing SAI Ethics necessitates

the integration of ethical principles into the design and development of AI algorithms and models. This means translating ethical principles, values, and guidelines into actionable and measurable practices or procedures. We need to define specific rules, standards, or protocols that guide the behavior and decision-making in ethical dilemmas or concrete situations [16, 17]. Moreover, SAI Ethics emphasizes the importance of continuous learning and adaptation. As AI technologies evolve and their societal impact unfolds, ethical standards and norms must evolve in tandem [18, 19]. This requires interdisciplinary research, ethical reflection, and stakeholder engagement to address emerging challenges and dilemmas.

4.2. Building a Computational Model of SAI Ethics

Ethical Principles are abstract rules intended for guiding ethical decision making and judgement. There are a variety of techniques used for technical implementation of ethical principles. In the previous literature of machine ethics, ethical principles are integrated into machines in a top-down, bottom-up, or hybrid architectures (see [20] for a survey). However, so far, no model seems to satisfy ethical judgement and decision making needs for an acceptable and responsible AI system. Approaches to encode principles into a format that computers can understand include logical reasoning, probabilistic reasoning, learning, optimisation, and case-based reasoning [21].

We argue that it is impossible to build a 'general ethical AI', i.e., a machine that is generally ethical, a machine that can reason and take ethical decisions in any domain and in every context. We believe that we need to concentrate on building *domain-based ethical machines*, i.e., machines that are able of ethical reasoning and decision making in any context and situation in a specific domain, which is, any way, still a very challenging task. Considering the purpose and the specific domain for which the AI system is developed, developers should consider codes of ethics and conduct of the domain (domain ethics, e.g. medical ethics) as a guiding framework. Furthermore, the key aspects of SAI, such as the collaborative and cooperative nature between human and machine, the human-centric approach, the mutual benefit, the adaptability and responsiveness of SAI, and the interdisciplinary perspective, should be taken in consideration in the design decisions to be taken by the developers.

To build a computational model of domain ethics to be integrated into the AI system; the ethical principles of the domain should be operationalized. The operationalization task should be carried out involving all stakeholders and domain ethical experts. Developers should also decide on the architecture to adopt for integrating the ethical principles. Being clear about which princi-

ple is being used will help designers to further specify what inputs are necessary for their application, which in turn will improve the ethical reasoning capabilities and explainability of how decisions have been made [22].

However, defining principles in an intentional manner so that they may be applied in a deductive manner, is often challenging and, in many cases, appears to be an impossible task. The issue lies in the gap between abstract, open-textured principles and tangible, concrete facts. The abstract principles should be operationalized by linking them to the facts. When ethical experts justify their conclusions in particular cases, they frequently connect ethical principles directly to the specific facts of those cases. Essentially, these established connections between ethical principles and relevant facts serve as operational (concrete) definitions of the principles. The experts operationalize the abstract principles by tying them directly to the factual context.

We are going to investigate, computationally, the possibility of operationalizing abstract ethical principles by inducing practical rules for ethical judgement and decision making in SAI systems from real-life interactions between human and machine in different domains [19, 23]. These rules evolve overtime through the interaction between human and machine which is an important aspect to SAI ethics. SAI recognizes the dynamic nature of human-AI interactions and the need for AI systems to adapt and respond to human preferences, values, and feedback overtime. To achieve this, we are going to consider different domains as case studies, collect and analyze a large set of domain ethics cases and build a computational model employing different operationalization techniques. Then, we are planning to carry out experiments to test our hypothesis that the computational model will accurately classify actions as ethical or unethical. The model will be developed using a foundational set of cases that will be collected for this purpose. The system performance will be evaluated using quantitative measures like precision and recall.

An important aspect, mentioned above, is the *model adaptability* overtime. In the context of SAI systems, human and machine (as agents) work as a team, collaborate and learn from each other, evolve together. The machine (as well as the human) will learn concrete ethical rules from interaction with humans, the machine will apply the previously learned ethical rules on concrete cases, will also revise and update the previously learned rules if needed. Here, it is important to emphasize the collaborative aspect of SAI in revising and correcting the ethical behavior overtime by both the human and the machine. In fact, this task is, in reality, a collaborative task, the machine will extract the case facts (the facts of the real-life case at hand), present them to the human, the human will provide an ethical judgment of the case at hand. Then the machine will learn a new rule and/or

revise a previously learned rule and present it to the human. Through a collaborative dialogue, The human can correct the ethical behavior of the machine, but also the machine can automatically demonstrate to the humans their errors in reasoning. In this way both will learn and improve their reasoning capabilities (mutual benefit). This adaptability aspect will be tested and evaluated in our experiments.

5. Conclusions and Future Work

In this work, we reported on ongoing work in the Work-Package 6.5 of the project FAIR. A model of ethical acceptability of SAI was outlined. Many legal issues raised by SAI systems were addressed. Currently, we are concentrating on SAI ethics operationalization. Next, we will work on the operationalization of legal aspects in SAI by the development of a framework for embedding the considerations of legal issues in SAI, then on realizing a computational model of legal reasoning for our SAI system to be ultimately integrated in the SAI system together with the ethical model.

By operationalizing SAI Ethics and legal issues, we can foster a collaborative and mutually beneficial relationship between humans and AI systems, promoting responsible and trustworthy AI development for the benefit of the society. This requires a multifaceted approach that integrates technical, organizational, regulatory, and societal perspectives.

A socio-technical approach to SAI systems development will be adopted which leads to an increased acceptability of these systems [24]. To capture the socio-technical complexity we are planning to adopt Multi-Agent Systems (MAS) for modelling the SAI system at hand [25]. The ethical and legal components in the system will be implemented as a MAS, which will act as an ethical and legal over-layer in the overall decision making process. A starting point might be the MAS prototype presented in [26, 27] for the ethical evaluation and monitoring of dialogue systems.

Finally, since a **human-centric approach** is central to SAI, transparency and explainability are key requirements for establishing trust in SAI systems which leads to acceptability. We would like to emphasize the the prominent role of computational logic in the development of the computational model of ethical and legal acceptability of SAI. Logic Programming (LP) has a great potential for developing such perspective ethical and legal SAI systems, as in fact logic rules are easily comprehensible by humans. Furthermore, LP is able to model causality, which is crucial for ethical and legal decision making [15].

Acknowledgments

This work was partially supported by the project FAIR - Future AI Research (PE00000013), under the NRRP MUR program funded by the NextGenerationEU.

References

- [1] J. C. R. Licklider, Man-computer symbiosis, *IRE Transactions on Human Factors in Electronics HFE-1* (1960) 4–11. doi:10.1109/THFE2.1960.4503259.
- [2] S. S. Grigsby, Artificial intelligence for advanced human-machine symbiosis, in: D. Schmorrow, C. Fidopiastis (Eds.), *Augmented Cognition: Intelligent Technologies*, volume 10915 of *Lecture Notes in Computer Science*, Springer, Cham, 2018. doi:10.1007/978-3-319-91470-1_22.
- [3] A. Carnevale, Condizione e struttura del nostro rapporto con le macchine. dieci proposizioni per una filosofia critica dell'intelligenza artificiale antropocentrica, in: S. Barone, et al. (Eds.), *L'uomo animale tecnologico*, Sciascia Editore, Caltanissetta-Rome, 2024. Invited chapter, accepted, in publication.
- [4] A. Lombardi, L'origine dell'io. il "mistero" dell'intelligenza da Darwin al riduzionismo contemporaneo, *Studium/Ricerca* 119 (2023) 651–688.
- [5] A. Carnevale, A. Lombardi, F. A. Lisi, Exploring ethical and conceptual foundations of human-centred symbiosis with artificial intelligence, in: G. Boella, et al. (Eds.), *Proceedings of the 2nd Workshop on Bias, Ethical AI, Explainability and the role of Logic and Logic Programming* co-located with the 22nd International Conference of the Italian Association for Artificial Intelligence (AI*IA 2023), volume 3615 of *CEUR Workshop Proceedings*, 2023, pp. 30–43. URL: <https://ceur-ws.org/Vol-3615/paper3.pdf>.
- [6] A. Carnevale, A. Lombardi, F. A. Lisi, A human-centred approach to symbiotic AI: Questioning the ethical and conceptual foundation, *Intelligenza Artificiale* (2024). Invited paper, in publication.
- [7] A. Carnevale, Assessing the impacts of symbiotic AI (SAI) on individual and societal well-being, in: H. Webb, et al. (Eds.), *AI Impact Assessment: methods and practices*, Oxford University Press, 2024. Invited chapter, accepted, in publication.
- [8] C. Falchi Delgado, M. T. Ferretti, A. Carnevale, Beyond one-size-fits-all: Precision medicine and novel technologies for sex and gender-inclusive covid-19 pandemic management, in: D. Cirillo, et al. (Eds.), *Innovating Health against Future Pandemics*, Elsevier, 2024. Invited chapter, accepted, in publication.
- [9] P. Marra, I. Galatola, Effectiveness as Threat to Constitutional Systems, Springer International Publishing, Cham, 2022, pp. 1–19. doi:10.1007/978-3-319-31739-7_142-1.
- [10] L. Pulito, Algoritmi predittivi e valutazione della pericolosità, *L'Ircocervo* (2024). Invited essay, submitted.
- [11] P. Marra, L. Pulito, A. Carnevale, F. Lisi, A. Lombardi, A. Dyoub, A procedural idea of decision-making in the context of symbiotic ai, in: *Proceedings of the 1st International Workshop on Designing and Building Hybrid Human-AI Systems*, co-located with 17th International Conference on Advanced Visual Interfaces (AVI 2024), Arenzano (Genoa), Italy, June 3rd, 2024., *CEUR Workshop Proceedings*, 2024. URL: <https://synergy.trx.li/ceur-ws/paper9.pdf>.
- [12] J. Morley, L. Kinsey, A. Elhalal, F. Garcia, M. Ziosi, L. Floridi, Operationalising AI ethics: barriers, enablers and next steps, *AI Soc.* 38 (2023) 411–423. URL: <https://doi.org/10.1007/s00146-021-01308-8>. doi:10.1007/s00146-021-01308-8.
- [13] J. Mökander, L. Floridi, Operationalising AI governance through ethics-based auditing: an industry case study, *AI Ethics* 3 (2023) 451–468. URL: <https://doi.org/10.1007/s43681-022-00171-7>. doi:10.1007/s43681-022-00171-7.
- [14] C. Novelli, F. Casolari, A. Rotolo, M. Taddeo, L. Floridi, AI risk assessment: A scenario-based, proportional methodology for the AI act, *Digit. Soc.* 3 (2024) 13. URL: <https://doi.org/10.1007/s44206-024-00095-1>. doi:10.1007/s44206-024-00095-1.
- [15] A. Dyoub, S. Costantini, F. A. Lisi, Logic programming and machine ethics, in: *Proceedings 36th International Conference on Logic Programming (Technical Communications), ICLP Technical Communications 2020, (Technical Communications) UNICAL, Rende (CS), Italy, 18-24th September 2020, volume 325 of EPTCS*, 2020, pp. 6–17. doi:10.4204/EPTCS.325.6.
- [16] A. Dyoub, S. Costantini, F. A. Lisi, Learning answer set programming rules for ethical machines, in: A. Casagrande, E. G. Omodeo (Eds.), *Proceedings of the 34th Italian Conference on Computational Logic*, Trieste, Italy, June 19-21, 2019, volume 2396 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2019, pp. 300–315. URL: <http://ceur-ws.org/Vol-2396/paper14.pdf>.
- [17] A. Dyoub, S. Costantini, F. A. Lisi, Towards an ILP application in machine ethics, in: *Inductive Logic Programming - 29th International Conference, ILP 2019, Plovdiv, Bulgaria, September 3-5, 2019, Proceedings*, volume 11770 of *Lecture Notes in Computer Science*, Springer, Netherlands, 2019, pp. 26–35. doi:10.1007/978-3-030-49210-6.
- [18] A. Dyoub, S. Costantini, I. Letteri, Care robots learning rules of ethical behavior under the supervision

- of an ethical teacher (short paper), in: P. Bruno, F. Calimeri, F. Cauteruccio, M. Maratea, G. Terracina, M. Vallati (Eds.), *Joint Proceedings of the 1st International Workshop on HYbrid Models for Coupling Deductive and Inductive ReAsoning (HYDRA 2022) and the 29th RCRA Workshop on Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion (RCRA 2022) co-located with the 16th International Conference on Logic Programming and Non-monotonic Reasoning (LPNMR 2022)*, Genova Nervi, Italy, September 5, 2022, volume 3281 of *CEUR Workshop Proceedings*, CEUR-WS.org, Germany, 2022, pp. 1–8. URL: <http://ceur-ws.org/Vol-3281/paper1.pdf>.
- [19] A. Dyoub, S. Costantini, F. A. Lisi, Learning domain ethical principles from interactions with users, *Digital Society 1 (2022)* 28. doi:10.1007/s44206-022-00026-y.
- [20] S. Tolmeijer, M. Kneer, C. Sarasua, M. Christen, A. Bernstein, Implementations in machine ethics: A survey, *ACM Computing Surveys* 53 (2020) 1–38. URL: <http://dx.doi.org/10.1145/3419633>. doi:10.1145/3419633.
- [21] S. J. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach*, Pearson Education Limited, 2016.
- [22] D. Leben, Normative principles for evaluating fairness in machine learning, in: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, AIES '20*, Association for Computing Machinery, New York, NY, USA, 2020, p. 86–92. URL: <https://doi.org/10.1145/3375627.3375808>. doi:10.1145/3375627.3375808.
- [23] A. Dyoub, S. Costantini, F. A. Lisi, I. Letteri, Logic-based machine learning for transparent ethical agents, in: F. Calimeri, S. Perri, E. Zumpano (Eds.), *Proceedings of the 35th Italian Conference on Computational Logic - CILC 2020*, Rende, Italy, October 13-15, 2020, volume 2710 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020, pp. 169–183. URL: <https://ceur-ws.org/Vol-2710/paper11.pdf>.
- [24] G. Baxter, I. Sommerville, Socio-technical systems: From design methods to systems engineering, *Interacting with Computers* 23 (2011) 4–17. doi:10.1016/j.intcom.2010.07.003.
- [25] K. Van Dam, I. Nikolic, Z. Lukszo, *Agent-Based Modelling of Socio-Technical Systems*, 2013. doi:10.1007/978-94-007-4933-7.
- [26] A. Dyoub, S. Costantini, I. Letteri, F. A. Lisi, A logic-based multi-agent system for ethical monitoring and evaluation of dialogues, in: A. Formisano, Y. A. Liu, B. Bogaerts, A. Brik, V. Dahl, C. Dodaro, P. Fodor, G. L. Pozzato, J. Vennekens, N. Zhou (Eds.), *Proceedings 37th International Conference on Logic Programming (Technical Communications)*, ICLP Technical Communications 2021, Porto (virtual event), 20-27th September 2021, volume 345 of *EPTCS*, 2021, pp. 182–188. doi:10.4204/EPTCS.345.32.
- [27] A. Dyoub, S. Costantini, F. A. Lisi, G. De Gasperis, Demo paper: Monitoring and evaluation of ethical behavior in dialog systems, in: Y. Demazeau, T. Holvoet, J. M. Corchado, S. Costantini (Eds.), *Advances in Practical Applications of Agents, Multi-Agent Systems, and Trustworthiness. The PAAMS Collection - 18th International Conference, PAAMS 2020*, L'Aquila, Italy, October 7-9, 2020, *Proceedings*, volume 12092 of *Lecture Notes in Computer Science*, Springer, UK, 2020, pp. 403–407. doi:10.1007/978-3-030-49778-1_35.