

# Automated Optimisation of Modern Software System Properties

Federica Sarro  
f.sarro@ucl.ac.uk  
University College London  
London, UK

## ABSTRACT

Realizing modern software systems poses new challenges to the software engineers: Users of applications running on limited capability devices still demand acceptable performance [2, 5, 13, 15]; users of systems relying on artificial intelligence to take decision (rightly) reclaim a fair treatment [4, 7, 12]; users of social networking systems expect to be protected against malicious behaviours [1]. Moreover, AI-enabled software systems are so energy-greedy that their usage is causing an alarming surge in energy consumption with a significant increase in CO<sub>2</sub> emissions [10].

Equipping software with appealing functionalities and minimising faults, is not enough if the emerging non-functional properties of these systems, such as fairness, safety and sustainability, are not taken into account. Mobile users will stop using an app if it is too slow or uses much bandwidth [5, 13]. Human bias can be transferred to various real-world systems relying on ML: Bias has been found in advertisement, recruitment, admission processes [3, 9, 19], among others, and human rights [16]. A growing number of malicious users use well-intentioned software platforms as a tool to attack the innocent users with whom they share the platform. Examples of such harmful acts are sadly too many to list; they include bullying, harassment, hate speech, misinformation, election interference, scamming and spamming. ChatGPT is an AI model able to answer a variety of questions, compose essays, have philosophical conversations, and even code or fix bugs [18]. However, all these come at a high cost: ChatGPT has been estimated to consume the equivalent of the electricity consumed by 175,000 people in Denmark per month.<sup>1</sup>

In this keynote, I will discuss the necessity to take these properties into account when realizing these type of systems, and the extent to which it is possible to automate their optimization. I will discuss existing solutions mainly based, but not limited to, multi-objective optimisation [5, 6, 8, 10, 14, 17]. In fact, we cannot expect that a software engineer, regardless of their level of expertise, would be able to manually find all opportunities for optimising these non-functional properties [11]. I will review research trends, presenting results from the SOLAR group and others.<sup>2</sup> I will also discuss some directions for future work and open-challenges towards achieving better, fairer, safer and greener software.

<sup>1</sup><https://towardsdatascience.com/chatgpts-electricity-consumption-7873483feac4>  
<sup>2</sup><https://solar.cs.ucl.ac.uk/>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).  
ICPE '23, April 15–19, 2023, Coimbra, Portugal  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0068-2/23/04.  
<https://doi.org/10.1145/3578244.3583739>

## CCS CONCEPTS

• **Software and its engineering** → **Software functional properties; Extra-functional properties; Search-based software engineering.**

## KEYWORDS

Software optimisation; functional properties; non-functional properties; search-based software engineering; genetic improvement; green software; green artificial intelligence

## ACM Reference Format:

Federica Sarro. 2023. Automated Optimisation of Modern Software System Properties. In *Proceedings of the 2023 ACM/SPEC International Conference on Performance Engineering (ICPE '23)*, April 15–19, 2023, Coimbra, Portugal. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3578244.3583739>

## BIO

Federica Sarro is a Professor of Software Engineering in the Department of Computer Science at UCL, where she is the Head of the Software Systems Engineering group and has established the SOLAR team within the CREST centre.



Her expertise lies at the intersection of software engineering, software optimisation, machine learning and evolutionary computation. She has carried out research on Machine Learning for Software Engineering and Software Engineering for Machine Learning with a focus on automatic software testing, repair, optimisation, software fairness, and predictive analytics for project management and quality assurance. She has published over 100 peer-reviewed scholarly articles, and given several invited talks at academic and industrial international events. She has led the international community in the development of predictive modelling for Software Engineering, and led scientific advances in the areas of Search-Based Software Engineering and Empirical Software Engineering.

Professor Sarro has obtained numerous awards and generous funding for her research. She has also worked with several companies including Meta, formerly known as Facebook, Google and Microsoft. In 2021, she was awarded the *Rising Star Award* by the IEEE Technical Community on Software Engineering in recognition of her “*excellence in Software Engineering research with scholarly and real-world impact*”.

Website: <http://www0.cs.ucl.ac.uk/staff/F.Sarro/>.

## ACKNOWLEDGMENTS

I am grateful to the members of my research group at UCL and the academic and industrial colleagues with whom I have had the pleasure to work so far. This work is supported by the ERC Advanced Grant no. 741278.

## REFERENCES

- [1] John Ahlgren, Kinga Bojarczuk, Sophia Drossopoulou, Inna Dvortsova, Johann George, Natalija Gucevska, Mark Harman, Maria Lomeli, Simon M. M. Lucas, Erik Meijer, Steve Omohundro, Rubmary Rojas, Silvia Sapora, and Norm Zhou. 2021. Facebook’s Cyber–Cyber and Cyber–Physical Digital Twins. In *Evaluation and Assessment in Software Engineering* (Trondheim, Norway) (EASE 2021). Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/3463274.3463275>
- [2] Afnan A. Al-Subaihini, Federica Sarro, Sue Black, Licia Capra, and Mark Harman. 2021. App Store Effects on Software Engineering Practices. *IEEE Transactions on Software Engineering* 47, 2 (2021), 300–319. <https://doi.org/10.1109/TSE.2019.2891715>
- [3] Peter J Bickel, Eugene A Hammel, and J William O’Connell. 1975. Sex bias in graduate admissions: Data from Berkeley. *Science* 187, 4175 (1975), 398–404.
- [4] Yuriy Brun and Alexandra Meliou. 2018. Software fairness. In *Proc. of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 754–759.
- [5] James Callan, Oliver Krauss, Justyna Petke, and Federica Sarro. 2022. How Do Android Developers Improve Non-Functional Properties of Software? *Empirical Softw. Engg.* 27, 5 (sep 2022), 42 pages. <https://doi.org/10.1007/s10664-022-10137-2>
- [6] Joymallya Chakraborty, Suvodeep Majumder, Zhe Yu, and Tim Menzies. 2020. Fairway: A Way to Build Fair ML Software. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (Virtual Event, USA) (ESEC/FSE 2020). Association for Computing Machinery, New York, NY, USA, 654–665. <https://doi.org/10.1145/3368089.3409697>
- [7] Zhenpeng Chen, Jie M. Zhang, Max Hort, Federica Sarro, and Mark Harman. 2022. Fairness Testing: A Comprehensive Survey and Analysis of Trends. <https://doi.org/10.48550/ARXIV.2207.10223>
- [8] Zhenpeng Chen, Jie M. Zhang, Federica Sarro, and Mark Harman. 2022. MAAT: A Novel Ensemble Approach to Addressing Fairness and Performance Bugs for Machine Learning Software. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (Singapore, Singapore) (ESEC/FSE 2022). Association for Computing Machinery, New York, NY, USA, 1122–1134. <https://doi.org/10.1145/3540250.3549093>
- [9] Jeffrey Dastin. 2018. Amazon scraps secret AI recruiting tool that showed bias against women. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- [10] Stefanos Georgiou, Maria Kechagia, Tushar Sharma, Federica Sarro, and Ying Zou. 2022. Green AI: Do Deep Learning Frameworks Have Different Costs?. In *44th IEEE/ACM International Conference on Software Engineering (ICSE)*. IEEE/ACM.
- [11] Mark Harman. 2022. Scaling Genetic Improvement and Automated Program Repair. In *2022 IEEE/ACM International Workshop on Automated Program Repair (APR)*. 1–7. <https://doi.org/10.1145/3524459.3527353>
- [12] Max Hort, Zhenpeng Chen, Jie M. Zhang, Federica Sarro, and Mark Harman. 2022. Bias Mitigation for Machine Learning Classifiers: A Comprehensive Survey. <https://doi.org/10.48550/ARXIV.2207.07068>
- [13] Max Hort, Maria Kechagia, Federica Sarro, and Mark Harman. 2021. A Survey of Performance Optimization for Mobile Applications. *IEEE Transactions on Software Engineering (TSE)* (2021).
- [14] Max Hort, Rebecca Moussa, and Federica Sarro. 2022. Multi-objective Search for Gender-fair and Semantically Correct Word Embeddings. In *Applied Soft Computing Journal*. IEEE.
- [15] William Martin, Federica Sarro, Yue Jia, Yuanyuan Zhang, and Mark Harman. 2017. A Survey of App Store Analysis for Software Engineering. *IEEE Transactions on Software Engineering* 43, 9 (2017), 817–847. <https://doi.org/10.1109/TSE.2016.2630689>
- [16] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [17] Rebecca Moussa, Giovanni Guizzo, and Federica Sarro. 2022. MEG: Multi-objective Ensemble Generation for Software Defect Prediction. In *16th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM 2022)*. ACM/IEEE.
- [18] Dominik Sobania, Martin Briesch, Carol Hanna, and Justyna Petke. 2023. An Analysis of the Automatic Bug Fixing Performance of ChatGPT. <https://doi.org/10.48550/ARXIV.2301.08653>
- [19] Jieyu Zhao, Yichao Zhou, Zeyu Li, Wei Wang, and Kai-Wei Chang. 2018. Learning Gender-Neutral Word Embeddings. *EMNLP* (2018), 4847–4853.