

Dr Francesco Leofante

last updated: October 20, 2026

Contact: f.leofante@imperial.ac.uk

Group webpage: <https://the-crumbs-lab.github.io>

Qualifications

2015 - 20 *PhD in Computer Science*. RWTH Aachen University (DE) & University of Genoa (IT).

2014 - 15 *MSc in Advanced Robotics*, Ecole Centrale de Nantes (FR).

2013 - 14 *MSc Robotics Engineering*, University of Genoa (IT).

2010 - 13 *BSc in Electronics Engineering*, University of Genoa (IT).

Career

2025 - present *Assistant Professor*. Department of Computing, Imperial College London (UK).

- Co-Director of the Centre for Explainable AI.

2022 - 2025 *Imperial College Research Fellow*. Department of Computing, Imperial College London (UK).

2020 - 22 *Research Associate*. Department of Computing, Imperial College London (UK).

2021 - present *AI consultant*. Imperial Consultants Ltd (UK).

Research grants

2023 *UK-Italy Trustworthy AI Visiting Researcher Programme*. Alan Turing Institute. Amount: **£4500**.

2022 *Imperial College Research Fellowship*. Imperial College London. Duration: 4 years. Amount: **~£250,000**.

2021 *TechCelerate accelerator grant*. Imperial College London. Amount: **£3000**.

2019 *AI*IA Mobility Grant*. Italian Association for Artificial Intelligence. Amount: **€2000**.

2018 *Travel grants*: AAAI'18, IJCAI-ECAI'18, IFM'18. Cumulative amount: **~£3000**.

Teaching and supervision

PhD level

2025 *Safe and Explainable AI*. PhD Course at the University of Genoa (12hrs).

2021 and 2022 *Introduction to AI Planning*. UKRI CDT on Safe and Trusted AI (4hrs, ~15 students).

MSc level

2027 *Causality and Explanation in AI*. Imperial College London (28hrs).

2024 - 2026 *Ethics, Fairness and Explanation in AI*. Imperial College London (12hrs, ~50 students).

2020 *Planning as Satisfiability*. University of Genoa (16hrs, ~15 students).

Undergraduate level

2025 *Discrete Maths, Logic and Reasoning*. Imperial College London (21hrs, ~270 students).

PhD supervision (2025-today) I am currently supervising:

- Mansi, AI controllability (started October 2025).
- Emanuele Albini, XAI for time series (co-supervised with Francesca Toni since September 2025).
- Andjela Todorovic, explainable GNNs (started October 2026).
- Vincent Zhang (co-supervised with Francesca Toni since October 2026)

Evidence of esteem

Awards

- 2026** IJCAI-ECAI 2026 Early Career Spotlight (corresponding invited paper [32]).
- 2024** Runner-up for the Pragnesh Jay Modi Best Paper Award at AAMAS'24 (paper [26]).
- 2023** Supporting Research Staff and Students Award. Imperial College London.
- 2018** 1st place in the Planning and Execution Competition for Logistics Robots in Simulation at ICAPS.
- 2018** Best presentation award at the PhD Symposium at IFM (paper [4]).

Professional memberships and fellowships

- 2026** Advance HE Fellow (FHEA).
- 2026** Member of ELLIS (European Laboratory for Learning and Intelligent Systems).
- 2024** DAAD (German Academic Exchange Service) AINet Fellowship.
- 2021** TechCelerate Fellow. Selected to join the Imperial TechCelerate accelerator programme.

Selected invited talks

- 2026** *What if...counterfactual explanations were to be deployed?* IJCAI Early Career Spotlight.
- 2026** *Explainable AI: the key to trustworthy cyber-defense.* London AI Summit.
- 2024** *Robustness Issues in Algorithmic Recourse.* Responsible AI Workshop, DERI Institute.

Selected invited seminars

- 2026** *Explaining AI through counterfactuals.* City St George's, University of London.
- 2025** *Towards counterfactual explanations for sequential decision-making.* Cardiff University.
- 2024** *Robust Algorithmic Recourse.* TU Dortmund and Bielefeld University.
- 2024** *Robustness (and Fairness) in Algorithmic Recourse.* Goldman Sachs.
- 2023** *Verification Techniques for Robust Explainable AI.* King's College London.
- 2023** *Robust Explainable AI.* Fondazione Bruno Kessler and University of Trento.

Tutorials at international venues

- 2023** *Robust Explainable AI: the Case of Counterfactual Explanations.* ECAI.
Website: <https://fraleo.github.io/ecai23-tutorial/>
- 2022** *SMT Solving and AI Planning.* EU ETN REMARO PhD School (Invited).
Website: <https://remaro.eu/index.php/remaro-fall-school-at-rwth-aachen-university/>
- 2018** *SMT Solving for AI Planning: Theory, Tools and Applications.* ICAPS.
Website: <https://ths.rwth-aachen.de/research/talks/smt4planning/>

Workshop organisation

- 2026** *Explainable AI and Computer Security (XAISEC).* Workshop at EuroS&P 2026.
Website: <https://xaisec-workshop.org>
- 2026** *Explainable Logic-Based Knowledge Representation & Explanations with Constraints and Satisfiability.* Workshop at FLOC 2026. Website: <https://sites.google.com/view/xlokrexcos2026/>

- 2026** *Explainable AI in Energy and Critical Infrastructure Systems*. Dagstuhl Seminar.
Website: <https://www.dagstuhl.de/seminars/seminar-calendar/seminar-details/26312>
- 2025** *Explainable AI in Energy and Critical Infrastructure Systems*. Bridge programme at AAAI 2025.
Website: <https://www.doc.ic.ac.uk/~fleofant/aaai25-xai-eci/index.html>
- 2022** *ExplAI 2022: Robustness in Explainable AI*. Imperial College London.
Website: <https://www.doc.ic.ac.uk/~fleofant/explainAI22/index.html>

Programme Committee membership AAAI'21,22,23,24,25,26,27}, AAMAS'22,23 (BlueSky),26}, ECAI'23,24,25} FAccT'24,25}, IJCAI'20,21,22,23,24,25,26 (SPC)}, KR'21,22,23,24,25,26}, NeurIPS'26.

Reviewer for journals Artificial Intelligence Journal, Minds and Machines, Logic Journal of the IGPL, Machine Learning, Information and Computation, Cognitive Systems Research.

Reviewer for national agencies Czech Academy of Sciences (CZ).

Jury member for international awards CLAIRE Rising Researcher Network Paper Highlight 2024.

Outreach and external engagement

- 2025** Energy Futures Podcast: AI for Critical Infrastructure Systems. URL: <https://www.youtube.com/watch?v=WmQkUtUUHH8>.
- 2025** Energy Futures Podcast: AI in Energy. URL: https://youtu.be/Vr6QLcz-L5U?si=mMAR_YuFJKhp3Zea.
- 2025** Doctoral Consortium @ AAAI 2025. Mentor for PhD students in AI.
- 2024** Robust Explainable AI. Outreach session at the European Conference on AI (ECAI 2024).
- 2024** Explainable AI (and other facts about AI research). Lecture to A-level students visiting Imperial.

Policy engagement

- 2025-2027** Member of Ofgem's AI Regulation Laboratory panel of experts.
- 2025-2026** AI Policy Fellowship @ Imperial. Mentor for senior civil servants enrolled in the programme.
- 2024** Contributor to the Parliamentary Office of Science and Technology research note "The future of energy security through AI and machine learning". DOI: <https://doi.org/10.58248/PN735>

Publications

Books

- [1] **Francesco Leofante** and Matthew Wicker. *Robust Explainable AI*. Springer Briefs in Intelligent Systems. Springer, 2025. ISBN: 978-3-031-89022-2. URL: <https://link.springer.com/book/9783031890215>.

Journals

- [7] L. Marzari, **F. Leofante**, F. Cicalese, and A. Farinelli. "Probabilistically Robust Counterfactual Explanations under Model Changes". In: *Artificial Intelligence* (2026).
- [6] E. Kyrimi et al. "Explainable AI: definition and attributes of a good explanation for health AI". In: *AI Ethics* (2025).
- [5] **F. Leofante**, A. Artelt, D. Eliades, A. Korre, F. Toni, and T. Miller. "Explainable AI, Energy and Critical Infrastructure Systems". In: *AI Magazine* (2025).

- [4] J. Jiang, **F. Leofante**, A. Rago, and F. Toni. "Interval Abstractions for Robust Counterfactual Explanations". In: *Artificial Intelligence* (2024).
- [3] **F. Leofante**. "OMTPlan: a Tool for Optimal Planning Modulo Theories". In: *Journal on Satisfiability, Boolean Modeling and Computation* (2023).
- [2] D. Guidotti, **F. Leofante**, A. Tacchella, and C. Castellini. "Improving Reliability of Myocontrol Using Formal Verification". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* (2019).
- [1] **F. Leofante**, E. Ábrahám, T. Niemueller, G. Lakemeyer, and A. Tacchella. "Integrated Synthesis and Execution of Optimal Plans for Multi-Robot Systems in Logistics". In: *Information Systems Frontiers* (2019).

Conferences

- [36] J. Jiang, **F. Leofante**, A. Rago, and F. Toni. "Synthesising Counterfactual Explanations via Label-Conditional Gaussian Mixture Variational Autoencoders". In: *Proc. of ICLR*. 2026 (to appear).
- [35] P. Kobiálka, A. Pferscher, **F. Leofante**, E. Ábrahám, S. Tapia Tarifa, and E. Broch Johnsen. "Attribution-based Explanations for Markov Decision Processes". In: *Proc. of IJCAI*. 2026.
- [34] Mansi, A. Kori, and **F. Leofante**. "Efficient bias mitigation in T2I diffusion models using Concept Graphs". In: *Proc. of NeurIPS*. 2026 (to appear).
- [33] R. Appachi Senthilkumar, **F. Leofante**, and V. Ganesh. "Enhancing the Robustness of Counterfactual Explanations via Adversarial Training". In: *Proc. of SAIV*. 2026.
- [32] **F. Leofante**. "What if...Counterfactual Explanations Were to Be Deployed?" In: *Proc. of IJCAI*. 2026.
- [31] G. Alfano, A. Gould, **F. Leofante**, A. Rago, and F. Toni. "Counterfactual Explanations under Model Multiplicity and their Use in Computational Argumentation". In: *Proc. of IJCAI*. 2025.
- [30] N. Gigante, **F. Leofante**, and A. Micheli. "Counterfactual Scenarios for Automated Planning". In: *Proc. of KR*. 2025.
- [29] J. Jiang, T. Bewley, S. Amoukou, **F. Leofante**, A. Rago, S. Mishra, and F. Toni. "Representation Consistency for Accurate and Coherent LLM Answer Aggregation". In: *Proc. of NeurIPS*. 2025.
- [28] J. Jiang, L. Marzari, A. Purohit, and **F. Leofante**. "RobustX: Robust Counterfactual Explanations Made Easy". In: *Proc. of IJCAI*. 2025.
- [27] P. Kobiálka, L. Gerlach, **F. Leofante**, E. Ábrahám, S. Tapia Tarifa, and E. Broch Johnsen. "Counterfactual Strategies for Markov Decision Processes". In: *Proc. of IJCAI*. 2025.
- [26] J. Jiang, **F. Leofante**, A. Rago, and F. Toni. "Recourse under Model Multiplicity via Argumentative Ensembling". In: *Proc. of AAMAS*. Runner-up for the Pragnesh Jay Modi Best Paper Award. 2024.
- [25] J. Jiang, **F. Leofante**, A. Rago, and F. Toni. "Robust Counterfactual Explanations in Machine Learning: A Survey". In: *Proc. of IJCAI*. 2024.
- [24] L. Marzari, **F. Leofante**, F. Cicalese, and A. Farinelli. "Rigorous Probabilistic Guarantees for Robust Counterfactual Explanations". In: *Proc. of ECAI*. The first two authors contributed equally. 2024.
- [23] **F. Leofante** and N. Potyka. "Promoting Counterfactual Robustness through Diversity". In: *Proc. of AAAI*. 2024.
- [22] **F. Leofante** et al. "Contestable AI needs Computational Argumentation". In: *Proc. of KR*. 2024.
- [21] J. Jiang, J. Lan, **F. Leofante**, A. Rago, and F. Toni. "Provably Robust and Plausible Counterfactual Explanations for Neural Networks via Robust Optimisation". In: *Proc. of ACML*. 2023.
- [20] J. Jiang, **F. Leofante**, A. Rago, and F. Toni. "Formalising the Robustness of Counterfactual Explanations for Neural Networks". In: *Proc. of AAAI*. The first two authors contributed equally. 2023.
- [19] P. Kouvaros, **F. Leofante**, B. Edwards, C. Chung, D. Margineantu, and A. Lomuscio. "Verification of Semantic Key Point Detection for Aircraft Pose Estimation". In: *Proc. of KR*. 2023.
- [18] **F. Leofante**, E. Botoeva, and V. Rajani. "Counterfactual Explanations and Model Multiplicity: a Relational Verification View". In: *Proc. of KR*. 2023.
- [17] **F. Leofante**, P. Henriksen, and A. Lomuscio. "Verification-friendly Networks: The Case for Parametric ReLUs". In: *Proc. of IJCNN*. The first two authors contributed equally. 2023.

- [16] **F. Leofante** and A. Lomuscio. "Robust Explanations for Human-Neural Multi-agent Systems with Formal Verification". In: *Proc. of EUMAS*. 2023.
- [15] **F. Leofante** and A. Lomuscio. "Towards Robust Contrastive Explanations for Human-Neural Multi-agent Systems". In: *Proc. of AAMAS*. 2023.
- [14] P. Henriksen, **F. Leofante**, and A. Lomuscio. "Repairing Misclassifications in Neural Networks Using Limited Data". In: *Proc. of SAC*. The first two authors contributed equally. 2022.
- [13] P. Kouvaros, T. Kyono, **F. Leofante**, A. Lomuscio, D. Margineantu, D. Osipychiev, and Y. Zheng. "Formal Analysis of Neural Network-based Systems in the Aircraft Domain". In: *Proc. of FM*. 2021.
- [12] D. Guidotti, **F. Leofante**, L. Pulina, and A. Tacchella. "Verification of Neural Networks: Enhancing Scalability through Pruning". In: *Proc. of ECAI*. 2020.
- [11] **F. Leofante**, E. Giunchiglia, E. Ábrahám, and A. Tacchella. "Optimal Planning Modulo Theories". In: *Proc. of IJCAI*. 2020.
- [10] A. Bit-Monnot, **F. Leofante**, L. Pulina, and A. Tacchella. "SMT-based Planning for Robots in Smart Factories". In: *Proc. of IEA/AIE*. 2019.
- [9] D. Guidotti, **F. Leofante**, L. Pulina, and A. Tacchella. "Verification and Repair of Neural Networks: A Progress Report on Convolutional Models". In: *Proc. of AI*IA*. 2019.
- [8] D. Guidotti, **F. Leofante**, A. Tacchella, and C. Castellini. "Repairing learned controllers with convex optimization: A case study". In: *Proc. of CPAIOR*. 2019.
- [7] **F. Leofante**, S. Schupp, E. Ábrahám, and A. Tacchella. "Engineering Controllers for Swarm Robotics via Reachability Analysis in Hybrid Systems". In: *Proc. of ECMS*. 2019.
- [6] **F. Leofante**. "Guaranteed Plans for Multi-Robot Systems via Optimization Modulo Theories". In: *Proc. of AAAI*. 2018.
- [5] **F. Leofante**. "Optimal Multi-robot Task Planning: From Synthesis to Execution (and Back)". In: *Proc. of IJCAI*. 2018.
- [4] **F. Leofante**, E. Ábrahám, and A. Tacchella. "Task Planning with OMT: An Application to Production Logistics". In: *Proc. of IFM*. 2018.
- [3] **F. Leofante**, E. Ábrahám, T. Niemueller, G. Lakemeyer, and A. Tacchella. "On the Synthesis of Guaranteed-Quality Plans for Robot Fleets in Logistics Scenarios via Optimization Modulo Theories". In: *Proc of IRI*. 2017.
- [2] **F. Leofante** and A. Tacchella. "Learning in Physical Domains: Mating Safety Requirements and Costly Sampling". In: *Proc. of AI*IA*. 2016.
- [1] **F. Leofante**, S. Viotto, E. Ábrahám, A. Tacchella, and N. Jansen. "Combining Static and Runtime Methods to Achieve Safe Standing-Up for Humanoid Robots". In: *Proc. of ISoLA*. 2016.