

JOÃO FONSECA

joaopfonseca.github.io
joaofonseca@inesc-id.pt ◇ joaofonseca@nyu.edu

POSITIONS

Postdoctoral Fellow, Distributed, Parallel and Secure Systems, INESC-ID <i>Advisors: Paolo Romano & Rodrigo Rodrigues</i>	Lisbon, Portugal Nov 2025 – Present
Visiting Scholar, Center for Responsible AI, New York University <i>Advisor: Julia Stoyanovich</i>	Remote Oct 2025 – Present
Postdoctoral Fellow, Center for Responsible AI, New York University <i>Advisor: Julia Stoyanovich</i>	New York, NY, USA Oct 2024 – Sep 2025
Invited Assistant Professor, NOVA University Lisbon <i>NOVA Information Management School</i>	Lisbon, Portugal Sep 2023 – Sep 2024
Visiting Scholar, Center for Responsible AI, New York University <i>Advisor: Julia Stoyanovich</i>	Remote Jul 2023 – Sep 2024
Research Fellow, NOVA University Lisbon <i>Advisor: Fernando Bação NOVA Information Management School</i>	Lisbon, Portugal Apr 2019 – Sep 2020

EDUCATION

NOVA University Lisbon	Lisbon, Portugal
Ph.D. in Information Management (NOVA Information Management School)	Sep 2020 – Oct 2023
M.S. in Information Management (NOVA Information Management School)	Sep 2017 – Jan 2019
M.S. in Management (NOVA School of Business and Economics)	Sep 2016 – Jan 2019
B.S. in Economics (NOVA School of Business and Economics)	Sep 2013 – Jun 2016
New York University	New York, NY, USA
Visiting PhD student / Research Intern (Tandon School of Engineering)	Mar 2023 – Jun 2023
Universidad de Granada	Granada, Spain
Exchange Semester (Facultad de Ciencias Económicas y Empresariales)	Sep 2015 – Feb 2016

PUBLICATIONS

Asterisks denote equal contribution or alphabetical ordering. Also see Google Scholar.

- [1] A. Bell*, J. **Fonseca***, C. Abrate, F. Bonchi, and J. Stoyanovich, “How much effort is enough? fairness in algorithmic recourse through the lens of substantive equality of opportunity,” in *Proceedings of the 5th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, ser. EAAMO ’25, New York, NY, USA: Association for Computing Machinery, 2025, pp. 170–184. [Online]. Available: <https://doi.org/10.1145/3757887.3763014>.
- [2] J. **Fonseca***, A. Bell*, and J. Stoyanovich, “SAFENUDGE: Safeguarding large language models in real-time with tunable safety-performance trade-offs,” in *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, ser. EMNLP ’25, Suzhou, China: Association for Computational Linguistics, Nov. 2025, pp. 19 955–19 969. [Online]. Available: <https://doi.org/10.18653/v1/2025.emnlp-main.1010>.
- [3] H. Hwang, A. Bell, J. **Fonseca**, V. Pliatsika, J. Stoyanovich, and S. E. Whang, “SHAP-based explanations are sensitive to feature representation,” in *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, ser. FAccT ’25, New York, NY, USA: Association for Computing Machinery, 2025, pp. 1588–1601. [Online]. Available: <https://doi.org/10.1145/3715275.3732105>.

[4] V. Pliatsika*, J. **Fonseca***, K. Akhynko, I. Shevchenko, and J. Stoyanovich, “Sharp: Explaining rankings and preferences with shapley values,” *Proc. VLDB Endow.*, vol. 18, no. 11, pp. 4131–4143, Jul. 2025. [Online]. Available: <https://doi.org/10.14778/3749646.3749682>.

[5] A. Bell*, J. **Fonseca***, and J. Stoyanovich, “The game of recourse: Simulating algorithmic recourse over time to improve its reliability and fairness,” in *Companion of the 2024 International Conference on Management of Data*, ser. SIGMOD ’24, Santiago, Chile: Association for Computing Machinery, 2024, pp. 464–467. [Online]. Available: <https://doi.org/10.1145/3626246.3654742>.

[6] **J. Fonseca**, “The role of synthetic data in improving supervised learning methods: The case of land use/land cover classification,” Ph.D. dissertation, Universidade NOVA de Lisboa (Portugal), 2023.

[7] **J. Fonseca** and F. Bacao, “Geometric smote for imbalanced datasets with nominal and continuous features,” *Expert Systems with Applications*, vol. 234, p. 121053, 2023. [Online]. Available: <https://doi.org/10.1016/j.eswa.2023.121053>.

[8] **J. Fonseca** and F. Bacao, “Improving active learning performance through the use of data augmentation,” *International Journal of Intelligent Systems*, vol. 2023, no. 1, p. 7941878, 2023. [Online]. Available: <https://doi.org/10.1155/2023/7941878>.

[9] **J. Fonseca** and F. Bacao, “Tabular and latent space synthetic data generation: A literature review,” *Journal of Big Data*, vol. 10, no. 1, p. 115, 2023. [Online]. Available: <https://doi.org/10.1186/s40537-023-00792-7>.

[10] **J. Fonseca***, A. Bell*, C. Abrate, F. Bonchi, and J. Stoyanovich, “Setting the right expectations: Algorithmic recourse over time,” in *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, ser. EAAMO ’23, Boston, MA, USA: Association for Computing Machinery, 2023. [Online]. Available: <https://doi.org/10.1145/3617694.3623251>.

[11] **J. Fonseca** and F. Bacao, “Research trends and applications of data augmentation algorithms,” *arXiv preprint arXiv:2207.08817*, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2207.08817>.

[12] **J. Fonseca**, G. Douzas, and F. Bacao, “Improving imbalanced land cover classification with k-means smote: Detecting and oversampling distinctive minority spectral signatures,” *Information*, vol. 12, no. 7, 2021. [Online]. Available: <https://doi.org/10.3390/info12070266>.

[13] **J. Fonseca**, G. Douzas, and F. Bacao, “Increasing the effectiveness of active learning: Introducing artificial data generation in active learning for land use/land cover classification,” *Remote Sensing*, vol. 13, no. 13, 2021. [Online]. Available: <https://doi.org/10.3390/rs13132619>.

[14] A. Crayton*, J. **Fonseca***, K. Mehra*, M. Ng*, J. Ross*, M. Sandoval-Castaneda*, and R. von Gnechten*, “Narratives and needs: Analyzing experiences of cyclone amphan using twitter discourse,” in *NeurIPS 2020 Workshop on Tackling Climate Change with Machine Learning*, 2020. [Online]. Available: <https://www.climatechange.ai/papers/neurips2020/42>.

[15] G. Douzas, F. Bacao, J. **Fonseca***, and M. Khudinyan*, “Imbalanced learning in land cover classification: Improving minority classes’ prediction accuracy using the geometric smote algorithm,” *Remote Sensing*, vol. 11, no. 24, 2019. [Online]. Available: <https://doi.org/10.3390/rs11243040>.

[16] **J. Fonseca**, “Harnessing big data to inform tourism destination management organizations,” M.S. thesis, Universidade NOVA de Lisboa (Portugal), 2018.

HONORS AND AWARDS

Cornell Tech – Runway Startup Postdoc Program (Rejected Offer)	2025
Catalyst NYC Accelerator – Pre-Seed Startup Accelerator	2025
Tech Venture Workshop – Non-Dilutive Startup Funding	2024
Best AI Track Paper Award – ACM EAAMO’23 Conference	2023
MIT Portugal Graduate Research Fellowship	2020
DSSG Summit 2020 Challenge – Hackathon Winner	2020

SELECTED INVITED TALKS

Promoting Responsible Innovation in Artificial Intelligence for Peace and Security exhibit, “Why do Large Language Models Hallucinate?”, United Nations, Office for Disarmament Affairs	May 7, 2025
AI Congressional Demo Speaker, U.S. Capitol Hill, Center for AI Policy	Feb 24, 2025
DSSG Summit 2020 – Presentation of the winning submission of the community challenge: “Keyboard layout optimization for ALS patients competition”	Oct 21, 2020
Solve for Good - DataFest 2020, “Amphan: Analyzing experiences of extreme weather events using online data”	Sep 17, 2020
Data Science for Hospitality and Tourism conference, “Unleashing the power of big data”, Panel member, Rectorate of NOVA University Lisbon	Dec 4, 2018
NOVA School of Business and Economics’ Inauguration Ceremony, “Data-driven planning for sustainable tourism in Portugal”	Sep 29, 2018

COMMUNITY SERVICE

- Reviewer: ICLR (2026, 2025), FAccT 2026

TEACHING

Computation II / Algorithms and Data Structures <i>Course coordinator NOVA Information Management School</i>	NOVA University Lisbon Feb 2024 – Jul 2024
Computation III: Object-oriented programming <i>Course coordinator NOVA Information Management School</i>	NOVA University Lisbon Sep 2023 – Jan 2024
Data Mining <i>Teaching Assistant NOVA Information Management School</i>	NOVA University Lisbon Sep 2019 – Jan 2024
Business Cases with Data Science <i>Teaching Assistant NOVA Information Management School</i>	NOVA University Lisbon Feb 2022 – Jul 2022
Introduction to Programming <i>Teaching Assistant NOVA School of Business and Economics</i>	NOVA University Lisbon Sep 2018 – Aug 2019

SUPERVISIONS

PhD Thesis	- Mafalda Sá Velho, “A New Paradigm in Explainable AI: Combining Large Language Models and XAI for Narrative Explanations in Black-Box Models”	2025 – PRES.
Masters Thesis	- Duarte Girão, “What must a player improve to increase its overall ranking?” - Reinaldo dos Santos Barros, “Improving relationship channels in the pharmaceutical industry with machine learning”	2024

LANGUAGES

Fluent in Portuguese, English and Spanish; Basic French.