

PASQUALE MINERVINI, CURRICULUM VITAE

Pasquale is a Lecturer in Natural Language Processing at the School of Informatics, University of Edinburgh; co-founder and CTO of the generative AI start-up [Miniml.AI](#); and an [ELLIS Scholar](#) – Edinburgh Unit. His research interests include NLP and ML, focusing on relational learning and learning from graph-structured data, solving knowledge-intensive tasks, hybrid neuro-symbolic models, compositional generalisation, and designing data-efficient and robust deep learning models. Pasquale routinely publishes in top-tier AI conferences and journals, receiving multiple awards (including one [Outstanding Paper Award](#) at ICLR 2021, one of the [most influential conferences in AI](#)), and delivered several tutorials in top-tier conferences (including five AAAI tutorials) – an updated list of his publications and tutorials is available on [his website](#). He is the Principal Investigator (PI) of the EU Horizon 2020 research grant *CLARIFY – Cancer Long Survivors Artificial Intelligence Follow Up*, the [Edinburgh Laboratory on Integrated Artificial Intelligence](#) (ELIAI) grant *Gradient-based Learning of Complex Latent Structures*, the Edinburgh-Huawei Joint Laboratory grant *Answering Complex Questions at Scale*, and leads several partnerships with other institutions and companies. In 2020, his team won two tracks out of three of the [Efficient Open-Domain Question Answering Challenge](#) at NeurIPS 2020; more recently, two of his PhD students ranked 1st in the [SemEval 2024 Task 2: Safe Biomedical Natural Language Inference for Clinical Trials](#) competition. Pasquale routinely collaborates with researchers across both academia and industry.

CONTACT DETAILS

OFFICE ADDRESS	Room 3.15, Informatics Forum, EH8 9AB Edinburgh, UK
E-MAIL ADDRESS	p.minervini@ed.ac.uk , p.minervini@miniml.ai
WEBSITE	https://neuralnoise.com
TWITTER/X	@pminervini

RESEARCH EXPERIENCE

09/2022 – Current	<i>Lecturer in Natural Language Processing</i> UNIVERSITY OF EDINBURGH, Edinburgh, UK Lecturer in Natural Language Processing at the Institute of Language, Cognition and Communication (ILCC) at the School of Informatics.
05/2024 – Current	<i>Co-founder and CTO</i> MINIML.AI , Edinburgh, UK The company focuses on developing reliable and trustworthy generative AI solutions for high-stakes domains and has ongoing projects with healthcare and financial brokerage companies.
05/2024 – Current	<i>Co-founder and CTO</i> MINIML.AI , Edinburgh, UK The company focuses on developing reliable and trustworthy generative AI solutions for high-stakes domains (e.g., finance and healthcare) and has ongoing projects with NHS and financial brokerage companies.
02/2019 – 09/2022	<i>Senior Research Fellow</i> UNIVERSITY COLLEGE LONDON, London, UK Senior Research Fellow in Statistical Natural Language Processing and Machine Learning in the UCL Natural Language Processing group. My position was fully funded by a three-year, $\approx$ 4.8 million EUR, Horizon 2020 grant that I manage as the PI on the UCL ends.
11/2020 – 12/2021	<i>Research Consultant</i> NEC LABORATORIES EUROPE GMBH, Heidelberg, Germany Research collaboration with Dr Mathias Niepert and his team.
10/2016 – 02/2019	<i>Research Fellow</i> UNIVERSITY COLLEGE LONDON, London, UK

Research Fellow in Statistical Natural Language Processing and Machine Learning in the UCL Machine Reading group. My position was funded by a Machine Reading grant from the Paul G. Allen Foundation, with Prof. Sebastian Riedel as PI.

- 12/2015 – 10/2016 *Postdoctoral Fellow*
INSIGHT CENTRE, NUI GALWAY, Galway, Ireland
Researcher in the area of knowledge discovery from both structured and unstructured data on the Web. The project is fully funded by and in close collaboration with Fujitsu Laboratories Ltd. and Fujitsu Ireland.
- 10/2015 – 12/2015 *Natural Language Processing (NLP) Engineer*
AYLIEN LTD., Dublin, Ireland
Research, development, and deployment of Deep Learning-based NLP models and systems: improved internal pre-existing NLP systems via neural architectures trained via distant supervision and labelled data and made them available via APIs to customers.
- 9/2014 – 9/2015 *Postdoctoral Fellow*
UNIVERSITÀ DEGLI STUDI DI BARI, Bari, Italy
Researcher for a research project titled “Methods and Techniques for Publishing and Mining in the Web of Data”.

EDUCATION

- MAY 2014 **PhD** in COMPUTER SCIENCE
Institution: **Università degli Studi di Bari**, Bari, Italy
Thesis Title: “Mining Methods for the Web of Data”
Advisor: Prof. Nicola Fanizzi Viva: May 26th, 2014
- FEBRUARY 2010 **Master’s Degree** in COMPUTER SCIENCE
Institution: **Università degli Studi di Bari**, Bari, Italy
Grade: **110/110, summa cum laude** (highest possible grade)
- FEBRUARY 2007 **Bachelor’s Degree** in COMPUTER SCIENCE
Institution: **Università degli Studi di Bari**, Bari, Italy
Grade: **110/110, summa cum laude** (highest possible grade)

RESEARCH FUNDING

- 2026 **TRANSpect**
Funding Body: **Medical Research Council (MRC)**
Value of Award: 527,928.79 GBP
Project: Beyond the Airlock: TRANSpect, Securing AI Transformer Models for Free-Text Safe Release in TREs
Duration: April 2026 – March 2027
Type of Grant: Research Funding
Role on the Grant: Co-I (PI: Arlene Casey)
- 2024 **Edinburgh-Huawei Joint Lab Grant**
Funding Body: **Huawei**
Value of Award: 296,919.43 GBP
Duration: September 2024 – September 2027
Type of Grant: Research Funding
Role on the Grant: PI
- 2024 **Generative AI Hub**

- Funding Body: **UKRI**
 Type of Grant: Research Funding
 Role on the Grant: Co-PI (PI: David Barber from UCL)
 More details available [at this link](#); more on EPSRC AI Hubs [here](#).
- 2024 **Big Ideas Accelerator (x2)**
 Funding Body: **University of Edinburgh**
 Type of Grant: Research Funding
 Role on the Grant: Co-PI
 I am Co-PI on two grants under this funding scheme, one on developing trustworthy and reliable foundation models for healthcare; and one on designing knowledge graphs for sociology research.
- 2022 **ELIAI Research Grant**
 Funding Body: **Edinburgh Laboratory for Integrated AI**
 Value of Award: 126,571.00 GBP
 Duration: January 2023 – January 2025
 Type of Grant: Research Funding
 Role on the Grant: PI (with Antonio Vergari and Edoardo Ponti from UoE)
- 2022 **Accenture Labs Grant**
 Funding Body: **Accenture Labs**
 Value of Award: 25,000 USD
 Duration: January 2023 – January 2024
 Type of Grant: Donation
 Role on the Grant: PI
- 2021 **Cisco Systems Research Grant**
 Funding Body: **Cisco Systems**
 Value of Award: 400,000 USD
 Duration: January 2022 – December 2024
 Type of Grant: Research Funding
 Role on the Grant: Co-PI (with Pontus Stenetorp from UCL)
- 2020 **Amazon Research Grant**
 Funding Body: **Amazon Research**
 Value of Award: 100,000 USD
 Duration: January 2021 – December 2021
 Type of Grant: Research Funding
 Role on the Grant: Co-PI (with Tim Rocktäschel from UCL)
- 2019 **Horizon 2020 Research Grant**
 Funding Body: **European Commission**
 Value of Award: 4,841,962.5 EUR
 Project: [CLARIFY – Cancer Long Survivors Artificial Intelligence Follow Up](#)
 Duration: January 2020 – December 2022
 Type of Grant: Research Funding
 Role on the Grant: Co-PI (with Pontus Stenetorp from UCL)
 I wrote the ML and NLP-related parts in the project proposal, and I am
 i) Financially, ethically, and legally managing the project on the UCL end,
 ii) Coordinating a WP, and iii) Contributing to several other WPs.
- 2016, 2018 **NVIDIA Academic Hardware Grant (2)**
 Funding Body: **NVIDIA Corporation**
 Value of Award: One NVIDIA Titan X GPU, One NVIDIA Titan Xp GPU
 Type of Grant: Hardware Grant
 Role on the Grant: PI

2011 – 2014 **Scholarship**

Funding Body: **Italian Ministry of Education, Universities and Research**

Value of Award: 46,000.00 EUR

Type of Grant: Scholarship

AWARDS

- SAC Highlights Award at [ACL 2025](#)
- Ranked 1st in the [SemEval 2024 Task 2: Safe Biomedical Natural Language Inference for Clinical Trials](#) competition. Award: 100.00 GBP.
- [Best Poster Award](#) at [ENLSP @ NeurIPS 2022](#). Award: 750.00 USD.
- [Outstanding Paper Award](#) at [ICLR 2021](#).
- [Ranked 1st](#) in two of the three tracks of the [Efficient Open-Domain Question Answering](#) challenge at NeurIPS 2020, together with collaborators from UCL and Facebook AI Research. Award: 3,000 USD in Google Cloud credits.
- Winner of the 4th LINKED DATA MINING CHALLENGE (Know@LOD 2016)
- Ranked 5th worldwide in the [Kaggle THE ALLEN AI SCIENCE Challenge](#), 2015
- [Best Research Paper Award](#) at the [19th International Conference on Knowledge Engineering and Knowledge Management](#) (EKAW 2014)
- Best Research Paper Award at the 10th International Workshop on Uncertainty Reasoning for the Semantic Web (URSW 2014)

PUBLICATIONS, TUTORIALS, AND PATENTS

PATENTS

- **Minervini**, Costabello, Muñoz, Nováček, Vandenbussche - Method and Apparatus for Completing a Knowledge Graph - US Patent Office, US Patent 11,341,417, Assignee: Fujitsu Ltd.

TUTORIALS

1. Lecue, Giannotti, Guidotti, **Minervini** - [On Explainable AI: From Theory to Motivation, Industrial Applications, XAI Coding & Engineering Practices](#). Tutorial for the Thirty-Seventh AAAI Conference on Artificial Intelligence (**AAAI** 2023)
2. Lecue, Giannotti, Guidotti, **Minervini** - [On Explainable AI: From Theory to Motivation, Industrial Applications, XAI Coding & Engineering Practices](#). Tutorial for the Thirty-Sixth AAAI Conference on Artificial Intelligence (**AAAI** 2022)
3. Lecue, **Minervini**, Giannotti, Guidotti - [On Explainable AI: From Theory to Motivation, Industrial Applications and Coding Practices](#). Tutorial for the Thirty-Fifth AAAI Conference on Artificial Intelligence (**AAAI** 2021)
4. Lecue, Gade, Geyik, Kenthapadi, Mithal, Taly, Guidotti, **Minervini** - [On Explainable AI: Foundations, Industrial Applications, Practical Challenges, and Lessons Learned](#). Tutorial for the Thirty-Fourth AAAI Conference on Artificial Intelligence (**AAAI** 2020)
5. Costabello, Lecue, Giannotti, Guidotti, Hitzler, **Minervini**, Sarker - [On Explainable AI: From Theory to Motivation, Applications and Limitations](#). Half-day (3.5 hours) Tutorial for the Thirty-Third AAAI Conference on Artificial Intelligence (**AAAI** 2019)
6. Guidotti, **Minervini**, Monreale, Rinzivillo - [Tutorial on Explainable Knowledge Discovery in Data Mining](#). Tutorial for the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (**ECML-PKDD** 2019)
7. Lecue, **Minervini** - [International Experts Tutorial on eXplainable AI](#). Artificial Intelligence Research Society, Seoul, Korea

Please note that while in other disciplines, conferences are generally just communication vehicles for presenting work, which is typically submitted to journals, Computer Science conferences publish proceedings of papers that are fully refereed publications and are often perceived to be of better quality and more prestigious than most journals. This is especially true for the top conferences in Artificial Intelligence (AAAI, IJCAI, UAI), Machine Learning (ICML, NeurIPS, ICLR, AAMAS), and Natural Language Processing (ACL, EMNLP, EACL, CoNLL), which are considered the best venue for publishing novel and ground-breaking results.

1. Leang, Zhao, Gema, Yang, Kwan, He, Li, **Minervini**, Giunchiglia, Cohen - PiCSAR: Probabilistic Confidence Selection And Ranking for Reasoning Chains. Findings of **ACL** 2026
2. Du, Hong, Kwan, Saxena, Titov, **Minervini**, Allaway - Analyzing LLM Instruction Optimization for Tabular Fact Verification. Findings of the Association for Computational Linguistics: **EACL** 2026
3. Gema, Hägele, Chen, Arditi, Goldman-Wetzler, Fraser-Taliente, Sleight, Petrini, Michael, Alex, **Minervini**, Chen, Benton, Perez - Inverse Scaling in Test-Time Compute. Transactions on Machine Learning Research (**TMLR** 2026)
4. Dileo, **Minervini**, Zignani, Gaito - Tensor Factorization for Temporal Knowledge Graph Forecasting. **Neurocomputing** 674 (2026)
5. Saxena, Suglia, **Minervini** - VLM-RobustBench: A Comprehensive Benchmark for Robustness of Vision-Language Models. CoRR abs/2603.06148 (2026)
6. Wani, Suglia, Saxena, Gema, Kwan, Barez, Bucarelli, Silvestri, **Minervini** - Same Answer, Different Representations: Hidden Instability in VLMs. CoRR abs/2602.06652 (2026)
7. Miller-Golub, Coil, Faber, Pietron, Zheng, **Minervini**, Corizzo - Rethinking the Harmonic Loss via Non-Euclidean Distance Layers. CoRR abs/2603.10225 (2026)
8. Lee, Yates, **Minervini**, Simpson - Process-Supervised Multi-Agent Reinforcement Learning for Reliable Clinical Reasoning. CoRR abs/2602.14160 (2026)
9. Kaddour, Patel, Dovonon, Richter, **Minervini**, Kusner - Agentic Uncertainty Reveals Agentic Overconfidence. CoRR abs/2602.06948 (2026)
10. van Krieken, **Minervini**, Ponti, Vergari - Neurosymbolic Diffusion Models. Thirty-Ninth Conference on Neural Information Processing Systems (**NeurIPS** 2025)
11. Wang, Yu, Ren, Zhang, Zhao, Saxena, Cheng, Wong, See, **Minervini**, Song, Steedman - MMLongBench: Benchmarking Long-Context Vision-Language Models Effectively and Thoroughly. Thirty-Ninth Conference on Neural Information Processing Systems (**NeurIPS** 2025, **spotlight**)
12. Richter, He, **Minervini**, Kusner - An Auditing Test to Detect Behavioral Shift in Language Models. 13th International Conference on Learning Representations (**ICLR** 2025)
13. Gregucci, Xiong, Hernández, Loconte, **Minervini**, Staab, Vergari - Is Complex Query Answering Really Complex? International Conference on Machine Learning (**ICML** 2025, **spotlight**, top 2.6% of accepted papers)
14. Zhu, de Souza, Shi, Yang, **Minervini**, D'Amour, Kusner - When Can Proxies Improve the Sample Complexity of Preference Learning? International Conference on Machine Learning (**ICML** 2025)
15. Hong, van Krieken, Ponti, Malkin, **Minervini** - Mixtures of In-Context Learners. 63rd Annual Meeting of the Association for Computational Linguistics (**ACL** 2025, SAC Highlights Award)
16. He, Wang, Xu, **Minervini**, Stenetorp, Rubinstein, Cohn - TuBA: Cross-Lingual Transferability of Backdoor Attacks in LLMs with Instruction Tuning. 63rd Annual Meeting of the Association for Computational Linguistics (**ACL** 2025)

17. Zheng, Chen, Fu, Wei, Zhang, **Minervini**, Ma, Gui, Zhang, Huang - GRADA: Graph-based Reranker against Adversarial Documents Attack. The 2025 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2025)
18. Arakelyan, **Minervini**, Verga, Lewis, Augenstein - FLARE: Faithful Logic-Aided Reasoning and Exploration. The 2025 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2025)
19. Gema, Jin, Abdulaal, Diethe, Teare, Alex, **Minervini**, Saseendran - DeCoRe: Decoding by Contrasting Retrieval Heads to Mitigate Hallucinations. The 2025 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2025)
20. Zhao, Devoto, Hong, Du, Gema, Wang, He, Wong, **Minervini** - Steering Knowledge Selection Behaviours in LLMs via SAE-Based Representation Engineering. The 2025 Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL** 2025)
21. Gema, Jun Leang, Hong, Devoto, Mancino, Saxena, He, Zhao, Du, Ghasemi Madani, Barale, McHardy, Harris, Kaddour, van Krieken, **Minervini** - Are We Done with MMLU? The 2025 Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL** 2025)
22. Luo, Gema, He, van Krieken, Lesci, **Minervini** - Self-Training Large Language Models for Tool-Use Without Demonstrations. The 2025 Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL** 2025, Findings)
23. Wójcik, Devoto, Pustelnik, **Minervini**, Scardapane - Adaptive Computation Modules: Granular Conditional Computation for Efficient Inference. The 39th AAAI Conference on Artificial Intelligence (**AAAI** 2025)
24. Ghazaryan, Arakelyan, **Minervini**, Augenstein - SynDARin: Synthesising Datasets for Automated Reasoning in Low-Resource Languages. The 30th International Conference on Computational Linguistics (**COLING** 2025)
25. Du, Saxena, Perez-Beltrachini, **Minervini**, Titov - Enhancing Long Document Long Form Summarisation with Self-Planning. IJCNLP-AAACL 2025
26. Saxena, **Minervini**, Keller - PosterSum: A Multimodal Benchmark for Scientific Poster Summarization. IJCNLP-AAACL 2025 (Findings)
27. Merra, Skracic, Malitesta, Golebiowski, **Minervini** - SmaLLEXT: 1st Workshop on Small and Efficient Large Language Models for Knowledge Extraction. 34th ACM International Conference on Information and Knowledge Management (**CIKM** 2025)
28. Saxena, Gema, **Minervini** - Lost in Time: Clock and Calendar Understanding Challenges in Multimodal LLMs. ICLR 2025 Workshop on Reasoning and Planning for Large Language Models
29. van Krieken, **Minervini**, Ponti, Vergari - Neurosymbolic Reasoning Shortcuts under the Independence Assumption. 19th International Conference on Neurosymbolic Learning and Reasoning (**NeSy** 2025)
30. Passerini, Gema, **Minervini**, Sayin, Tentori - Fostering Effective Hybrid Human-LLM Reasoning and Decision Making. Frontiers in Artificial Intelligence 7, 1464690 (2025)
31. Ghasemi Madani, Gema, Sarti, Zhao, **Minervini**, Passerini - Noiser: Bounded Input Perturbations for Attributing Large Language Models. Conference on Language Modeling (**CoLM** 2025)
32. Devoto, Alvetreti, Pomponi, Di Lorenzo, **Minervini**, Scardapane - Adaptive Layer and Token Selection for Efficient Fine-Tuning of Vision Transformers. **Neurocomputing** (2025)
33. Murphy, Rizvi, Haussmann, Nie, Liu, Gema, **Minervini** - An Analysis of Decoding Methods for LLM-based Agents for Faithful Multi-Hop Question Answering. CoRR abs/2503.23415 (2025)

34. Kwan, Leang, Vougiouklis, Pan, Valentino, **Minervini** - OpenSIR: Open-Ended Self-Improving Reasoner. CoRR abs/2511.00602 (2025)
35. Sundrani, Tudisco, **Minervini** - Low-Rank Compression of Language Models via Differentiable Rank Selection. CoRR abs/2512.13733 (2025)
36. Godey, Devoto, Zhao, Scardapane, **Minervini**, de la Clergerie, Sagot - Q-Filters: Leveraging QK Geometry for Efficient KV Cache Compression. CoRR abs/2503.02812 (2025)
37. Arakelyan, Hambardzumyan, Papikyan, **Minervini**, Gordo, Augenstein, Markosyan - With Great Backbones Comes Great Adversarial Transferability. CoRR abs/2501.12275 (2025)
38. Zhao, Chen, Inan, Kessler, Wang, Wutschitz, Yang, Zhang, **Minervini**, Sim - Learning GUI Grounding with Spatial Reasoning from Visual Feedback. CoRR abs/2509.21552 (2025)
39. Szatkowski, Bedkowski, Devoto, Dubinski, **Minervini**, Piorczynski, Scardapane, Wójcik - Universal Properties of Activation Sparsity in Modern Large Language Models. CoRR abs/2509.00454 (2025)
40. Sayin, Baris Schlicht, Vo Hong, Allievi, Staiano, **Minervini**, Passerini - MedSyn: Enhancing Diagnostics with Human-AI Collaboration. CoRR abs/2506.14774 (2025)
41. Grivas, Loconte, van Krieken, Nawrot, Zhao, Wielewski, **Minervini**, Ponti, Vergari - Fast and Expressive Multi-Token Prediction with Probabilistic Circuits. CoRR abs/2511.11346 (2025)
42. Luo, Recharadt, Sun, Nejad, Yáñez, Yilmaz, Lee, Cohen, Borghesani, Pashkov, Marinazzo, Nicholas, Salatiello, Sucholutsky, **Minervini**, Razavi, Rocca, Yusifov, Okalova, Gu, Ferianc, Khona, Patil, Lee, Mata, Myers, Bizley, Musslick, Bilgin, Niso, Ales, Gaebler, Murty, Loued-Khenissi, Behler, Hall, Dafflon, Bao, Love - Large language models surpass human experts in predicting neuroscience results. **Nature Human Behaviour** (2024)
43. Scardapane, Baiocchi, Devoto, Marsocci, **Minervini**, Pomponi - Conditional computation in neural networks: principles and research trends. **Intelligenza Artificiale** 18(1): 175–190 (2024)
44. Devoto, Zhao, Scardapane, **Minervini** - A Simple and Effective L_2 Norm-Based Strategy for KV Cache Compression. The 2024 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2024 Main Conference; Oral, top 9% of accepted papers)
45. Stacey, **Minervini**, Dubossarsky, Camburu, Rei - Atomic Inference for NLI with Generated Facts as Atoms. The 2024 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2024 Main Conference)
46. Zhou, Nie, Guo, Wei, Zhang, **Minervini**, Ma, Gui, Zhang, Huang - Unveiling and Consulting Core Experts in Retrieval-Augmented MoE-based LLMs. The 2024 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2024 Main Conference)
47. Zhao, Qu, Staniszewski, Tworkowski, Liu, Miłoś, Wu, **Minervini** - Analysing The Impact of Sequence Composition on Language Model Pre-Training. The 62nd Annual Meeting of the Association for Computational Linguistics (**ACL** 2024 Main Conference; Oral, top 8% of accepted papers)
48. van Krieken, **Minervini**, Ponti, Vergari - On the Independence Assumption in Neurosymbolic Learning. International Conference on Machine Learning (**ICML** 2024)
49. Solano, Sanni, Camburu, **Minervini** - SparseFit: Few-shot Prompting with Sparse Fine-tuning for Jointly Generating Predictions and Natural Language Explanations. The 62nd Annual Meeting of the Association for Computational Linguistics (**ACL** 2024 Main Conference)
50. He, Wu, Camburu, **Minervini**, Stenetorp - Using Natural Language Explanations to Improve Robustness of In-context Learning. The 62nd Annual Meeting of the Association for Computational Linguistics (**ACL** 2024 Main Conference)

51. Wang, **Minervini**, Ponti - Probing the Emergence of Cross-lingual Alignment During Large Language Model Pre-Training. The 62nd Annual Meeting of the Association for Computational Linguistics (**ACL 2024 Findings**)
52. Armengol-Estapé, Rocha, Woodruff, **Minervini**, O’Boyle - Forklift: An Extensible Neural Lifter. Conference on Language Modeling (**CoLM 2024**, 28% acceptance rate)
53. Yu, He, **Minervini**, Pan - Evaluating the Adversarial Robustness of Retrieval-Based In-Context Learning for Large Language Models. Conference on Language Modeling (**CoLM 2024**, 28% acceptance rate)
54. Tyukin, Dovonon, Kaddour, **Minervini** - Attention Is All You Need But You Don’t Need All Of It For Inference of Large Language Models. ICML 2024 Workshop on Theoretical Foundations of Foundation Model (**TF2M @ ICML 2024**)
55. Richter, Agrawal, He, **Minervini**, Kusner - An Auditing Test to Detect Behavioral Shift in Language Models. ICML 2024 Workshop on Foundation Models in the Wild (**FM-Wild @ ICML 2024**)
56. Gema, Hong, **Minervini**, Daines, Alex - Edinburgh Clinical NLP at SemEval-2024 Task 2: Fine-tune your model unless you have access to GPT-4. 7th ClinicalNLP@NAACL 2024 (**ClinicalNLP 2024**)
57. Gema, Lee, **Minervini**, Daines, Simpson, Alex - Edinburgh Clinical NLP at MEDIQA-CORR 2024: Guiding Large Language Models with Hints. 7th ClinicalNLP@NAACL 2024 (**ClinicalNLP 2024**)
58. Sayin, **Minervini**, Staiano, Passerini - Can LLMs Correct Physicians, Yet? Investigating Effective Interaction Methods in the Medical Domain. 7th ClinicalNLP@NAACL 2024 (**ClinicalNLP 2024**)
59. Gema, Daines, **Minervini**, Alex - Parameter-Efficient Fine-Tuning of LLaMA for the Clinical Domain. 7th ClinicalNLP@NAACL 2024 (**ClinicalNLP 2024**)
60. **Minervini**, Franceschi, Niepert - Adaptive Perturbation-Based Gradient Estimation for Discrete Latent Variable Models. 37th AAAI Conference on Artificial Intelligence (**AAAI 2023**, 19.6% acceptance rate)
61. Arakelyan*, **Minervini***, Daza, Cochez, Augenstein - Adapting Neural Link Predictors for Data-Efficient Complex Query Answering. 37th Conference on Neural Information Processing Systems (**NeurIPS 2023**, 26% acceptance rate)
62. Kaddour, Key, Nawrot, **Minervini**, Kusner - No Train No Gain: Revisiting Efficient Training Algorithms For Transformer-based Language Models. 37th Conference on Neural Information Processing Systems (**NeurIPS 2023**, 26% acceptance rate)
63. Andresel, Tran, Domokos, **Minervini**, Stepanova - Combining Inductive and Deductive Reasoning for Query Answering over Incomplete Knowledge Graphs. 32nd ACM International Conference on Information and Knowledge Management (**CIKM 2023**)
64. Zhou, Iacobacci, **Minervini** - XQA-DST: Multi-Domain and Multi-Lingual Dialogue State Tracking. 17th Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2023**)
65. Janik, Torrente, Costabello, Calvo, Walsh, Camps, Mohamed, Ortega, Nováček, Massutí, **Minervini**, Garcia Campelo, Del Barco, Bosch-Barrera, Menasalvas, Timilsina, Provencio - Machine Learning-Assisted Recurrence Prediction for Patients With Early-Stage Non-Small-Cell Lung Cancer. JCO Clinical Cancer Informatics 2023:7
66. Timilsina, Fey, Buosi, Janik, Costabello, Carcereny, Rodriguez Abreu, Cobo, López Castro, Bernabé, **Minervini**, Torrente, Provencio, Nováček - Synergy between imputed genetic pathway and clinical information for predicting recurrence in early-stage non-small cell lung cancer. Journal of Biomedical Informatics 144, ISSN 1532-0464
67. Timilsina, Buosi, Janik, **Minervini**, Costabello, Torrente, Provencio, Calvo, Camps, Ortega, Massutí, Garcia Campelo, del Barco Bosch-Barrera, Nováček - Learning Survival Models for

- Relapse Prediction in Early-Stage Lung Cancer Patients. International Joint Conference on Neural Networks (**IJCNN** 2023)
68. Cochez, Alivanistos, Arakelyan, Berrendorf, Daza, Galkin, Minervini, Niepert, Ren - Approximate Answering of Graph Queries. Chapter in *Compendium of Neurosymbolic Artificial Intelligence*, IOS Press, 2023
 69. Chen, Mishra, Franceschi, **Minervini**, Stenetorp, Riedel - ReFactor GNNs: Revisiting Factorisation-based Models from a Message-Passing Perspective. 36th Conference on Neural Information Processing Systems (**NeurIPS** 2022)
 70. Stacey, **Minervini**, Dubossarsky, Rei - Logical Reasoning with Span Predictions: Span-level Logical Atoms for Interpretable and Robust NLI Models. 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2022)
 71. Wu, Zhao, Hu, **Minervini**, Stenetorp, Riedel - An Efficient Memory-Augmented Transformer for Knowledge-Intensive NLP Tasks. 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2022) and Efficient Natural Language and Speech Processing Workshop (**ENLSP @ NeurIPS** 2022, [best poster award](#))
 72. Amin*, **Minervini***, Chang, Stenetorp, Neumann - MedDistant19: Towards an Accurate Benchmark for Broad-Coverage Biomedical Relation Extraction. 29th International Conference on Computational Linguistics (**COLING** 2022)
 73. **Minervini***, Arakelyan*, Daza*, Cochez - Complex Query Answering with Neural Link Predictors (Extended Abstract). 31st International Joint Conference on Artificial Intelligence (**IJCAI** 2022)
 74. Morris, **Minervini**, Blunsom - Learning Proof Path Selection Policies in Neural Theorem Proving. 16th International Workshop on Neural-Symbolic Learning and Reasoning (**NeSy @ IJCLR** 2022), and 4th Conference on Automated Knowledge Base Construction (**AKBC** 2022)
 75. Wren, **Minervini**, Franceschi, Zantedeschi - Learning Discrete Directed Acyclic Graphs via Backpropagation. Workshop on Causality for Real-world Impact (**CML4Impact @ NeurIPS** 2022), and Workshop on Neuro Causal and Symbolic AI (**nCSI @ NeurIPS** 2022)
 76. Arakelyan*, Daza*, **Minervini***, Cochez - Complex Query Answering with Neural Link Predictors. 9th International Conference on Learning Representations (**ICLR** 2021, **oral presentation** 2% acceptance rate, overall 29% – [Outstanding Paper Award](#))
 77. Niepert, **Minervini**, Franceschi - Implicit MLE: Backpropagating Through Discrete Exponential Family Distributions. 35th Conference on Neural Information Processing Systems (**NeurIPS** 2021)
 78. Jiang, **Minervini**, Jiang, Rocktäschel - Grid-to-Graph: Flexible Spatial Relational Inductive Biases for Reinforcement Learning. 20th International Conference on Autonomous Agents and Multiagent Systems (**AAMAS** 2021, 24.8% acceptance rate)
 79. Lewis, Wu, Liu, **Minervini**, Küttler, Piktus, Stenetorp, Riedel - PAQ: 65 Million Probably-Asked Questions and What You Can Do With Them. Transactions of the Association for Computational Linguistics (**TACL** 2021)
 80. Wu, **Minervini**, Stenetorp, Riedel - Training Adaptive Computation for Open-Domain Question Answering with Computational Constraints. The Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (**ACL-IJCNLP** 2021, 14.9% acceptance rate)
 81. **Minervini**, Riedel, Stenetorp, Grefenstette, Rocktäschel - Learning Reasoning Strategies in End-to-End Differentiable Proving - Chapter in *Neuro-Symbolic Artificial Intelligence: The State of the Art*, IOS Press, 2022
 82. Mohamed, Walsh, Timilsina, Torrente, Franco, Provencio, Janik, Costabello, **Minervini**, Stenetorp, Nováček - On Predicting Recurrence in Early Stage Non-small Cell Lung Cancer. American Medical Informatics Association Annual Symposium (**AMIA** 2021)

83. de Vassimon Manela, Errington, Fisher, van Breugel, **Minervini** - Stereotype and Skew: Quantifying Gender Bias in Pre-trained and Fine-tuned Language Models. 16th Conference of the European Chapter of the Association for Computational Linguistics (**EACL** 2021)
84. Dobrowolska, Vergari, **Minervini** - Neural Concept Formation in Knowledge Graphs. 3rd International Conference on Automated Knowledge Base Construction (**AKBC** 2021)
85. Chen, **Minervini**, Stenetorp, Riedel - Relation Prediction as an Auxiliary Training Objective for Improving Multi-Relational Graph Representations. 3rd International Conference on Automated Knowledge Base Construction (**AKBC** 2021)
86. Torrente, Franco, Calvo, Collazo Lorduy, Menasalvas, Vidal, Sousa, Pimentao, Novacek, **Minervini**, Fey, Costabello, Pocs, Provencio - P08.01 Building Personalized Follow-Up Care Through AI by Bringing the Lung Cancer Patient, Data Scientist and Oncologist Together. Journal of Thoracic Oncology, International Association for the Study of Lung Cancer 16(10):S991-S992, October 2021
87. Betz, Niepert, **Minervini**, Stuckenschmidt - Backpropagating through Markov Logic Networks. 15th International Workshop on Neural-Symbolic Learning and Reasoning (**NeSy** 2021)
88. Chauhan, Gupta, **Minervini** - A Probabilistic Framework for Knowledge Graph Data Augmentation. Data-Centric AI Workshop at NeurIPS (**DCAI** 2021)
89. **Minervini**, Riedel, Stenetorp, Grefenstette, Rocktäschel - Learning Reasoning Strategies in End-to-End Differentiable Proving - 37th International Conference on Machine Learning (**ICML** 2020, 21.8% acceptance rate)
90. **Minervini***, Bošnjak*, Rocktäschel, Riedel, Grefenstette - Differentiable Reasoning on Large Knowledge Bases and Natural Language - 34th AAAI Conference on Artificial Intelligence (**AAAI** 2020, **oral presentation** 4.5% acceptance rate; overall 20.6%)
91. Camburu, Shillingford, **Minervini**, Lukasiewicz, Blunsom - Make Up Your Mind! Adversarial Generation of Inconsistent Natural Language Explanations. 58th Annual Meeting of the Association for Computational Linguistics (**ACL** 2020, 17.6% acceptance rate)
92. Wu, **Minervini**, Stenetorp, Riedel - Don't Read Too Much Into It: Adaptive Computation for Open-Domain Question Answering. The 2020 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2020, 22.4% acceptance rate)
93. Stacey, **Minervini**, Dubossarsky, Riedel, Rocktäschel - Gone At Last: Removing the Hypothesis-Only Bias in Natural Language Inference via Ensemble Adversarial Training. The 2020 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2020, 22.4% acceptance rate)
94. Welbl, **Minervini**, Bartolo, Stenetorp, Riedel - Undersensitivity in Neural Reading Comprehension. The 2020 Conference on Empirical Methods in Natural Language Processing (**EMNLP** 2020, Findings)
95. Min, Boyd-Graber, Alberti, Chen, Choi, Collins, Guu, Hajishirzi, Lee, Palomaki, Raffel, Roberts, Kwiatkowski, Lewis, Wu, Küttler, Liu, **Minervini**, Stenetorp, Riedel, Yang, Seo, Izacard, Petroni, Hosseini, De Cao, Grave, Yamada, Shimaoka, Suzuki, Miyawaki, Sato, Takahashi, Suzuki, Fajcik, Docekal, Ondrej, Smrz, Cheng, Shen, Liu, He, Chen, Gao, Oguz, Chen, Karpukhin, Peshterliev, Okhonko, Schlichtkrull, Gupta, Mehdad, Yih - NeurIPS 2020 EfficientQA Competition: Systems, Analyses and Lessons Learned. NeurIPS 2020 Competition and Demonstration Track (**NeurIPS** 2020)
96. **Minervini**, Bošnjak, Rocktäschel, Riedel, Grefenstette - Differentiable Reasoning on Large Knowledge Bases and Natural Language - Chapter in *Knowledge Graphs for eXplainable Artificial Intelligence: Foundations, Applications and Challenges*, IOS Press, 2020
97. Bianchi, Rossiello, Costabello, Palmonari, **Minervini** - Knowledge Graph Embeddings and Explainable AI - Chapter in *Knowledge Graphs for eXplainable Artificial Intelligence: Foundations, Applications and Challenges*, IOS Press, 2020

98. Dobrowolska, Vergari, **Minervini** - Neural Concept Formation in Knowledge Graphs - Knowledge Representation & Reasoning Meets Machine Learning Workshop at NeurIPS (**KR2ML** 2020)
99. Camburu, Shillingford, **Minervini**, Lukasiewicz, Blunsom - Make Up Your Mind! Adversarial Generation of Inconsistent Natural Language Explanations - NeurIPS 2019 Workshop on Safety and Robustness in Decision Making (**SafeRobust@NeurIPS**, 2019)
100. Weber, **Minervini**, Münchmeyer, Leser, Rocktäschel - NLProlog: Reasoning with Weak Unification for Question Answering in Natural Language - 57th Annual Meeting of the Association for Computational Linguistics (**ACL** 2019, 22.7% acceptance rate)
101. Muñoz, **Minervini**, Nickles - Embedding Cardinality Constraints in Neural Link Predictors - 34th ACM/SIGAPP Symposium on Applied Computing (**ACM SAC** 2019, 25% acceptance rate)
102. Cowen-Rivers, **Minervini**, Riedel, Rocktäschel, Wang, Bošnjak - Neural Variational Inference for Estimating Knowledge Graph Embedding Uncertainty - 14th International Workshop on Neural-Symbolic Learning and Reasoning (**NeSy@IJCAI** 2019)
103. **Minervini**, Riedel - Adversarially Regularising Neural NLI Models to Integrate Logical Background Knowledge - SIGNLL Conference on Computational Natural Language Learning (**CoNLL** 2018, 20.65% acceptance rate)
104. **Minervini**, Bošnjak, Campero, Rocktäschel, Grefenstette, Riedel - Neural Theorem Proving on Natural Language - International Conference on Probabilistic Programming (**PROBPROG** 2018)
105. **Minervini**, Bošnjak, Rocktäschel, Riedel - Towards Neural Theorem Proving at Scale - Workshop on Neural Abstract Machines & Program Induction (**NAMPI** 2018)
106. Mitchell, **Minervini**, Stenetorp, Riedel - Extrapolation in NLP - Workshop on Generalization in the Age of Deep Learning (**NAACL** 2018)
107. Weissenborn, **Minervini**, Dettmers, Augenstein, Welbl, Rocktäschel, Bošnjak, Mitchell, Demeester, Stenetorp, Riedel - Jack the Reader – A Machine Reading Framework - Annual Meeting of the Association for Computational Linguistics (**ACL** 2018), System Demonstrations
108. Dettmers, **Minervini**, Stenetorp, Riedel - Convolutional 2D Knowledge Graph Embeddings - 31st AAAI Conference on Artificial Intelligence (**AAAI** 2018, 24.6% acceptance rate)
109. **Minervini**, Tresp, d’Amato, Fanizzi - Adaptive Knowledge Propagation in Web Ontologies - ACM Transactions on the Web (**TWEB** 2018)
110. **Minervini**, Demeester, Rocktäschel, Riedel - Adversarial Sets for Regularising Neural Link Predictors - 33rd Conference on Uncertainty in Artificial Intelligence (**UAI** 2017)
111. **Minervini**, Costabello, Muñoz, Nováček, Vandenbussche - Regularizing Knowledge Graph Embeddings via Equivalence and Inversion Axioms - European Conference on Machine Learning & Principles and Practice of Knowledge Discovery in Databases (**ECML-PKDD** 2017) (27% acceptance rate)
112. **Minervini**, d’Amato, Fanizzi - Efficient Energy-Based Embedding Models for Link Prediction in Knowledge Graphs - Journal on Intelligent Information Systems (**JiIS** 2016), Recent Advances in Mining Patterns from Complex Data, ISSN 1573-7675
113. **Minervini**, d’Amato, Fanizzi, Tresp - Discovering Similarity and Dissimilarity Relations for Knowledge Propagation in Web Ontologies - Journal on Data Semantics (**JoDS** 2016), ISSN 1861-2040
114. **Minervini**, d’Amato, Fanizzi - Leveraging the Schema in Latent Factor Models for Knowledge Graph Completion - Proceedings of the ACM Symposium on Applied Computing - Semantic Web Track (**ACM SAC** 2016), Pisa, Italy (24% acceptance rate)

115. Yumusak, Muñoz, **Minervini**, Dogdu, Kodaz - A Hybrid Method for Rating Prediction Using Linked Data Features and Text Reviews - (**KNOW@LOD/CoDeS@ESWC** 2016)
116. **Minervini**, d'Amato, Fanizzi, Esposito - Scalable Learning of Entity and Predicate Embeddings for Knowledge Graph Completion - 14th IEEE International Conference on Machine Learning and Applications, (**ICMLA** 2015), ISBN 978-1-5090-0287-0
117. **Minervini**, d'Amato, Fanizzi, Esposito - Efficient Learning of Entity and Predicate Embeddings for Link Prediction in Knowledge Graphs - Proceedings of the 11th International Workshop on Uncertainty Reasoning for the Semantic Web (**URSW** 2015)
118. **Minervini**, d'Amato, Fanizzi, Esposito - A Gaussian Process Model for Knowledge Propagation in Web Ontologies - IEEE International Conference on Data Mining (**ICDM** 2014), ISBN 978-1-4799-4302-9 (19% acceptance rate)
119. **Minervini**, d'Amato, Fanizzi, Esposito - Adaptive Knowledge Propagation in Web Ontologies - Proceedings of the 19th International Conference on Knowledge Engineering and Knowledge Management (**EKAW** 2014), ISBN 978-3-319-13703-2, Linköping, Sweden (**best research paper award**)
120. **Minervini**, d'Amato, Fanizzi, Esposito - Graph-Based Regularization for Transductive Class-Membership Prediction. In: Uncertainty Reasoning for the Semantic Web III - ISWC International Workshops, **URSW** 2011-2013, Revised Selected Papers - Springer, ISBN: 978-3-319-13412-3
121. **Minervini**, d'Amato, Fanizzi, Esposito - Learning Probabilistic Description Logic Concepts Under Alternative Assumptions on Incompleteness. In: Uncertainty Reasoning for the Semantic Web III - ISWC International Workshops, **URSW** 2011-2013, Revised Selected Papers - Springer, ISBN: 978-3-319-13412-3
122. **Minervini**, d'Amato, Fanizzi, Tresp - Learning to Propagate Knowledge in Web Ontologies. - Proceedings of the 10th International Workshop on Uncertainty Reasoning for the Semantic Web (**URSW** 2014, **best research paper award**)
123. **Minervini**, Fanizzi, d'Amato, Esposito - Rank Prediction for Semantically Annotated Resources - Proceedings of the ACM Symposium on Applied Computing - Semantic Web Track (**ACM SAC** 2013), ISBN 978-1-4503-1656-9 (24% acceptance rate)
124. **Minervini**, d'Amato, Fanizzi, Esposito - Transductive Inference for Class-Membership Propagation in Web Ontologies - The Semantic Web: Semantics and Big Data (**ESWC** 2013), ISBN 978-3-642-38287-1 (26% acceptance rate)
125. **Minervini**, d'Amato, Fanizzi - A Graph Regularization Based Approach to Transductive Class-Membership Prediction - Proceedings of the 8th International Workshop on Uncertainty Reasoning for the Semantic Web (**URSW** 2012)
126. Fanizzi, d'Amato, Esposito, **Minervini** - Numeric Prediction on OWL Knowledge Bases through Terminological Regression Trees - International Journal on Semantic Computing (**IJSC**) 2012
127. **Minervini**, d'Amato, Fanizzi - Learning Terminological Bayesian Classifiers: A Comparison of Alternative Approaches to Dealing with Unknown Concept Memberships - Proceedings of the 9th Italian Convention on Computational Logic (**CILC** 2012)
128. **Minervini**, d'Amato, Fanizzi - Learning Terminological Naive Bayesian Classifiers under Different Assumptions on Missing Knowledge. Proceedings of the 7th International Workshop on Uncertainty Reasoning for the Semantic Web (**URSW** 2011), CEUR Workshop Proceedings vol. 778 ISSN 1613-0073, pg. 63-74
129. **Minervini**, d'Amato, Fanizzi - Learning Probabilistic Description Logic Concepts Under Different Assumptions on Missing Knowledge. Proceedings of the ACM Symposium on Applied Computing - Semantic Web Track (**ACM SAC** 2012), ISBN 978-1-4503-0857-1 (26% acceptance rate)

130. Calefato, Lanubile, **Minervini** - Can Real-Time Machine Translation Overcome Language Barriers in Distributed Requirements Engineering? - 5th IEEE International Conference on Global Software Engineering (**ICGSE** 2010) ISBN 978-1-4244-7619-0
131. **Minervini** - Apertium goes SOA: an efficient and scalable service based on the Apertium rule-based machine translation platform - Proceedings of the First International Workshop on Free/Open-Source Rule-Based Machine Translation (**FreeRBMT** 2009), ISBN-13: 978-8-46-136188-5, pg. 59-65