

Curriculum Vitae

Felice Antonio Merra

April 11, 2023

Personal Data

mobile (DE): +49 0176 83 27 5356
mobile (IT): +39 340 42 13 560
e-mail: merrafelice@gmail.com
website: <https://merrafelice.github.io/>

Social Accounts:

Twitter:

Personal Account – > @merrafelice

Research Divulcation Account – > @AdversRecSys

LinkedIn: Felice Antonio Merra

GitHub: merrafelice

YouTube: Felice Antonio Merra

Current Position

Applied Scientist II. Amazon, Amazon Search, Berlin (DE).

Brief Description of Academic and Industrial Scientific Activities

My research activities mainly focus on the application of machine learning techniques in information retrieval. During the PhD, I devoted my attention to recommender systems (RSs) with a particular focus on the robustness of modern ML recommender models affected by adversarial threats. After having assessed the state-of-the-art of AML techniques in recommendation domain, I have investigated three main areas of study: (i) the robustness of recommender models when affected by hand-engineered shilling attacks, (ii) the formal study of the effects of adversarial training strategies on the beyond-accuracy effects of recommenders, i.e., bias disparity, fairness, novelty, and (iii) the proposal of adversarial attacks and defenses against multimedia retrieval models from image to music. The main research contributions are reported in my PhD thesis available at the following link.

My industrial research activities are focused on investigating novel algorithmic approaches for improving the relevance of Amazon Search system. Being involved in the search engine, I work withing Information Retrieval (IR) and Natural Language Processing (NLP) research areas. My activities as an applied scientist span from identify, explore, prototype and develop ML solutions for matching, retrieval, and ranking stages. Within my experience as an applied scientist, I have been involved in multiple projects.

In particular, I have developed a an extreme multi-label classification (XMLC) model to solve the filter selection problem whose has been presented at WWW2021. I have developed multiple solutions for the improvement of the search engine on users queries containing a negation intent. In particular, in this project, I have been the scientific point of contact and I have worked with multiple people in different locations to implement and test the proposed solution within a real setting. I have also focused my attention to document understanding research activities working as the tech leader of the project and mentoring a junior scientist.

My academic and industrial activities have required, and require, a constant connection with researchers within other research teams in different location (EMEA, USA, and Japan). I am active in publishing articles about my industrial research contributions inside internal conferences (and working also for external) with the goal to be always connected with the research community. I serve also as reviewers or member of program committee in several venues and journals, e.g., RecSys, TRS, or AAAI.

1 Education

Ph.D. in Computer Science and Automation Engineering. February 2022. Thesis Topic: *Adversarial Machine Learning in Recommender Systems.*. **Final Score:** Full marks with honors.

Master's Degree in Computer Science Engineering. October 2018. Thesis title: *Privacy As A Service: automatic anoyimization of microdataset for urban analytic.* In this work, a novel initial to data-quality preservation in Open Data has been proposed. **Final Score:** Full marks with honors.

Bachelor's Degree in Computer Science and Automation Engineering. July 2016. Thesis title: *Information Flow Processing: Complex Event Processing in Health-care Systems.* **Final Score:** Full marks with honors.

2 Awards

Best Short Paper Runner Up Award at the 30th ACM International Conference on Information and Knowledge Management, 1-5 November 2021, for the [12] article.

Best Paper Award at the 3rd Workshop on Adversarial Learning Methods for Machine Learning and Data Mining @ KDD 2021 sponsored by MIT-IBM Watson AI Lab for the [5] article.

3 Internship

Applied Science Intern at Amazon.com. July-September 2020. *Amazon Search and Personalization* research team.

4 Honors

- **Subject Expert:** Algorithms and Data Structure in Java *2020*
- **Poliba Ph.D. Fellowship:** Ph.D. supported by Politecnico di Bari *2018-2021*

- **Poliba Scholarship Award for Master Thesis Abroad:** Research Visiting Period at the Knowledge Media Institute (UK) supported by Politecnico di Bari 2018

5 Research Interests

My research activity starts in 2018, soon after my graduation at Polytechnic University of Bari (PoliBa). After a 4 month period as a Visiting Research at Knowledge Media Institute in UK under the supervision of Prof. Enrico Motta working on the privacy protection algorithms in demographic reports, I pursued a Ph.D. on Recommender Systems at PoliBa. During the second year of my PhD, I have been an Intern at Amazon.com working on recommendation algorithms for widgets selection. Soon after the completion of the Ph.D., I moved back to Amazon.com where I am working as an applied scientists in Amazon Search.

5.1 PhD Dissertation Abstract

Recommender systems are ubiquitous. Our digital lives are influenced by their use when, for instance, we select the news to read, the product to buy, the friend to connect, and the movie to watch. While enormous academic research efforts have been mainly focused on getting high-quality recommendations to reach the maximum customers' satisfaction, little effort has been devoted to studying the integrity and security of these systems. Is there an underlying relationship between the characteristics of the historical user-item interactions and the efficacy of injection of false users/feedback strategies against collaborative models? Can public semantic data be used to perform attacks more potent in raising the recommendability of victim items? Can a malicious user (i.e., the adversary) poison or evade the image data of visual recommenders with adversarial perturbed product images? What is a possible defensive solution to reduce the effectiveness of test-time adversarial attacks? Is the family of model-based recommenders more vulnerable to multi-step gradient-based adversarial perturbations? Furthermore, is the adversarial training robustification still effective in the last scenario? Is this training defense influencing the beyond-accuracy and bias performance?

My dissertation intended to pave the way towards more robust recommender systems, beginning with understanding how a model can be made more robust, the cost of robustness in terms of recommendation quality, and the adversarial risks of modern recommenders. My thesis, getting inspiration from the literature on the security of collaborative models against the insertion of hand-engineered fake profiles and the recent advances of adversarial machine learning methods in other research areas like computer vision, contributed to several directions: (i) the proposal of a practical framework to interpret the impact of data characteristics on the robustness of collaborative recommenders, (ii) the design of powerful attack strategies using publicly available semantic data, (iii) the identification of severe adversarial vulnerabilities of visual-based recommender models where adversaries can break the recommendation integrity by pushing products to the highest recommendation positions with a simple and human-imperceptible perturbation of products' images, (iv) the design of a novel defense method to protect visual recommenders against test-time adversarial attacks, (v) the proposal of robust adversarial perturbation methods capable of completely breaking the accuracy of matrix factorization recommenders, and (vi) a formal study that examines the effects of adversarial training in reducing the recommendation quality of state-of-the-art model-based recommenders.

6 International Collaboration

Within my research activities, I have worked and currently work, with the following international research institute:

- Amazon Science and Engineering teams spreading all over the world.
- Consorzio Interuniversitario per il Calcolo Automatico dell'Italia Nord Orientale (Cineca), Bologna, Italy Accepted Grant for the use of HPC Resources
- Escuela Politécnica Superior, Universidad Autónoma de Madrid (dr. Alejandro Bellogín), Madrid, Spain
- Università del Salento and Politecnico di Torino for the project *FLET4.0 – FLEet managemenT optimization through I4.0 enabled smart maintenance*, as Project Research Scientist
- Knowledge Media institute - The Open University (prof. Enrico Motta, dr. A. Antonini), Milton Keynes, United Kingdom

7 Tools Experience

[Programming Languages] I have strong experience in Python programming, with a daily use of TensorFlow and PyTorch. I have also worked on projects with other program languages such as C++, C#, Java and JavaScript. I consider my coding skills really important for my research path and career.

[Amazon SageMaker] I am a consumer of multiple Amazon SageMaker services, from notebooks to IMR clusters used for interacting with Big Data via PySpark, from pre-processing, training and inference jobs to Step and Lambda Functions. I have also contributed in the design of services used for large scaled products.

[Large Scale Experiments] During my Ph.D. activities, I have matured experience in performing large-scale experiments on Super Computing resources. In particular, I have been performing experiments on Marconi100 ¹, a 32 PFlops accelerated cluster based on IBM Power9 architecture.

[Multi-GPU Training for Feature Extraction on Music Content] In the GitHub project, i.e., link to the project, I have implemented in TensorFlow v2, a convolutional neural network, with multi-gpu training, for the extraction of latent features from the mel-spectrogram generated from music tracks.

8 International Activities

In the follow, I present the attended conferences, the activities as reviewer/sub-reviewer, and attended summer schools schools.

- Program Committee:
 - RecSys (Log and Short Paper), RecSys LBR, WDCS@NeurIPS, ECIR, JAIR, IEEE Signal Processing, AMLC
- Summer Schools:
 - RecSys 2019 Summer School, The ACM Summer School on Recommender Systems, September 9-13, 2019 Goteborg, Sweden
 - ESSIR 2019, The 12th European Summer School in Information Retrieval, July 15-19, 2019 Milan, Italy

¹<https://www.hpc.cineca.it/hardware/marconi100>

9 GitHub Projects

An updated set of repositories related to my research activities and published papers is available at: <https://github.com/merrafelice>.

[**Elliot**] I have actively contributed in the development of an **open-source library** for the realization of extensive and reproducible experiments on recommendation models. In particular, the framework implements the complete pipeline from the data preparation, to the training and evaluation of state-of-the-art recommendation models (more than 50 models), passing through a set of possible optimization search strategies (e.g., Bayesian optimization). The public daily-maintained repository is available at <https://github.com/sisinflab/elliott>. In particular, I have actively contributed into the integration of deep learning, adversarial learning, and visual,-based recommendations models published into a RecSys 2021 demonstration paper [5].

10 Teaching Activities

Data Structure and Algorithm in Java. I started my teaching activity in the academic year 2018/2019, 2019/20, 2020/2021 as Course Assistant for the bachelor degree course: *Data Structure and Algorithm in Java*.

11 Selected Publications

I list here four major works where I am the main author.

- **How Dataset Characteristics Affect the Robustness of Collaborative Recommendation Models.** ACM SIGIR 2020. In this long paper, I have worked on proposing a framework to understand whether recommendation dataset characteristics (e.g., rating variance and sparsity) can impact the performance of shilling attacks against pure collaborative recommendation models.
- **A Study of Defensive Methods to Protect Visual Recommendation Against Adversarial Manipulation of Images.** ACM SIGIR 2021. In this work, I have identified, for the first time in the recsys community, the vulnerabilities of state-of-the-art visual-based recommendation models when adversarial samples are uploaded by sellers on e-commerce platforms. The essence of this work has been the proposal of an overall framework that has taken into account SOTA adversarial techniques used to manipulate products' images together with possible defense approaches. Within this line of work, I had the chance to get experience with many visual-recommender, Computer Vision techniques for image processing (and attack), as well as, a deep understanding of visual-based limitation when the content is noised. Note that this line of research is connected to other minor publications where I studied other dimensions like biases, human perception, and the effects of CNNs. Currently, I am working on the proposal of a U-net-based Autoencoder to denoise adversarial images by preserving high-recommendation performance.
- **A Formal Analysis of Recommendation Quality of Adversarially-trained Recommenders.** ACM CIKM 2021. In this article, awarded with the MIT-IBM Watson AI Lab best paper award, I have opened the black-box of adversarial regularization in model-based recommenders studying the gradient-based characteristics of this technique with respect to its connection to possible amplification of popularity bias. In this case, I had the chance to understand at a very deep level how a recommender model learns users' preferences.

- **Search Filter Ranking with Language-Aware Label Embeddings.** ACM WWW 2022. In this work, we frame the filter selection task as an extreme multi-label classification (XMLC) problem based on historical interaction with eCommerce sites. We learn from customers' clicks and purchases which subset of filters is most relevant to their queries treating the relevant/not-relevant signal as binary labels.

12 Publications

The authors of the papers are sorted in alphabetical order. I am the **main author** of all the listed publications related to the **adversarial machine learning** topic.

- [1] Vito Walter Anelli, Alejandro Bellogin, Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. Msap: Multi-step adversarial perturbations on recommender systems embeddings. In *The 34th International FLAIRS Conference*, pages 1–6. The Florida AI Research Society (FLAIRS), AAAI Press, May 2021.
- [2] Vito Walter Anelli, Alejandro Bellogín, Antonio Ferrara, Daniele Malitesta, Felice Antonio Merra, Claudio Pomo, Francesco M. Donini, Eugenio Di Sciascio, and Tommaso Di Noia. The challenging reproducibility task in recommender systems research between traditional and deep learning models. In *SEBD*, volume 3194 of *CEUR Workshop Proceedings*, pages 514–521. CEUR-WS.org, 2022.
- [3] Vito Walter Anelli, Alejandro Bellogín, Antonio Ferrara, Daniele Malitesta, Felice Antonio Merra, Claudio Pomo, Francesco Maria Donini, and Tommaso Di Noia. Elliot: a comprehensive and rigorous framework for reproducible recommender systems evaluation. In *Proceedings of the 44th International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2021, Virtual Event, Canada*. ACM, July 2021.
- [4] Vito Walter Anelli, Alejandro Bellogín, Antonio Ferrara, Daniele Malitesta, Felice Antonio Merra, Claudio Pomo, Francesco Maria Donini, and Tommaso Di Noia. How to perform reproducible experiments in the elliot recommendation framework: data processing, model selection, and performance evaluation. In *IIR*, CEUR Workshop Proceedings. CEUR-WS.org, 2021.
- [5] Vito Walter Anelli, Alejandro Bellogín, Antonio Ferrara, Daniele Malitesta, Felice Antonio Merra, Claudio Pomo, Francesco Maria Donini, and Tommaso Di Noia. V-elliot: Design, evaluate and tune visual recommender systems. In *RecSys 2021: Fifteenth ACM Conference on Recommender Systems (RecSys '21), September 27-October 1, 2021, Amsterdam, Netherlands*. ACM, 2021.
- [6] Vito Walter Anelli, Alejandro Bellogín, Antonio Ferrara, Daniele Malitesta, Felice Antonio Merra, Claudio Pomo, Francesco Maria Donini, and Tommaso Di Noia. V-elliot: Speeding up visual recommendation via a gpu-powered data input pipeline. In *NVIDIA GPC 2022*, 2022.
- [7] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta, and Felice Antonio Merra. A study of defensive methods to protect visual recommendation against adversarial manipulation of images. In *The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, page 10. ACM, July 2021.

-
- [8] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. Adversarial learning for recommendation. In *Advances in Information Retrieval - 43rd European Conference on IR Research, ECIR 2021*, Lecture Notes in Computer Science. Springer, 2021.
- [9] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. *Adversarial Recommender Systems: Attack, Defense, and Advances*. to appear in the Third Edition of Recommender Systems Handbook, Springer, third edition, 2021. to appear in the Third Edition of Recommender Systems Handbook.
- [10] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. Understanding the effects of adversarial personalized ranking optimization method on recommendation quality. In *3rd Workshop on Adversarial Learning Methods for Machine Learning and Data Mining @ KDD 2021 (virtual workshop)*. Online, 2021.
- [11] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. Adversarial learning for recommendation: Applications for security and generative tasks - concept to code. In *RecSys 2020: Fourteenth ACM Conference on Recommender Systems, Virtual Event, Brazil, September 22-26, 2020*, pages 738–741. ACM, 2020.
- [12] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. A formal analysis of recommendation quality of adversarially-trained recommenders. In *CIKM*. ACM, 2021.
- [13] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. Understanding the effects of adversarial personalized ranking optimization method on recommendation quality. *3rd Workshop on Adversarial Learning Methods for Machine Learning and Data Mining co-located with KDD 2021*, 2021.
- [14] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Felice Antonio Merra, Giuseppe Acciani, and Eugenio Di Sciascio. Knowledge-enhanced shilling attacks for recommendation. In *Proceedings of the 28th Italian Symposium on Advanced Database Systems, Villasimius, Sud Sardegna, Italy (virtual due to Covid-19 pandemic), June 21-24, 2020*, volume 2646 of *CEUR Workshop Proceedings*, pages 310–317. CEUR-WS.org, 2020.
- [15] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, and Felice Antonio Merra. Sasha: Semantic-aware shilling attacks on recommender systems exploiting knowledge graphs. In *The Semantic Web - 17th International Conference, ESWC 2020, Heraklion, Crete, Greece, May 31-June 4, 2020, Proceedings*, volume 12123 of *Lecture Notes in Computer Science*, pages 307–323. Springer, 2020.
- [16] Vito Walter Anelli, Tommaso Di Noia, Daniele Malitesta, and Felice Antonio Merra. Assessing perceptual and recommendation mutation of adversarially-poisoned visual recommenders. *The 1st Workshop on Dataset Curation and Security co-located with the 34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada (Virtual Event)*., 2020.
- [17] Vito Walter Anelli, Tommaso Di Noia, Daniele Malitesta, and Felice Antonio Merra. Assessing perceptual and recommendation mutation of adversarially-poisoned visual recommenders (short paper). In *DP@AI*IA*, volume 2776 of *CEUR Workshop Proceedings*, pages 49–56. CEUR-WS.org, 2020.

-
- [18] Vito Walter Anelli, Tommaso Di Noia, and Felice Antonio Merra. The idiosyncratic effects of adversarial training on bias in personalized recommendation learning. In *RecSys 2021: Fifteenth ACM Conference on Recommender Systems (RecSys '21)*, September 27–October 1, 2021, Amsterdam, Netherlands. ACM, 2021.
- [19] Vito Walter Anelli, Tommaso Di Noia, Eugenio Di Sciascio, Daniele Malitesta Deldjoo, and Felice Antonio Merra. Adversarial attacks against visual recommendation: an investigation on the influence of items’ popularity. In *Proceedings of the Second Workshop on Online Misinformation- and Harm-Aware Recommender Systems co-located with 15th ACM Conference on Recommender Systems (RecSys 2021) Virtual Event*, Amsterdam, The Netherlands, October 2, 2021.
- [20] Giuseppe De Candia, Tommaso Di Noia, Eugenio Di Sciascio, and Felice Antonio Merra. Amflp: Adversarial matrix factorization-based link predictor in social graphs. In *SEBD 2021: The 29th Italian Symposium on Advanced Database Systems, September 5-9, 2021, Pizzo Calabro (VV), Italy*. CEUR Workshop Proceedings, sept 2021.
- [21] Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta, and Felice Antonio Merra. A study on the relative importance of convolutional neural networks in visually-aware recommender systems. In *CVPRW-CVFAD 2021 :The 4th CVPR Workshop on Computer Vision for Fashion, Art, and Design*, pages 1–4. CVPR Proceedings, June 2021.
- [22] Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta Deldjoo, and Felice Antonio Merra. Leveraging content-style item representation for visual recommendation. In *Proceeding of the 44th European Conference on Information Retrieval (ECIR 2022)*, 10-14 April 2022, Stavanger, Norway.
- [23] Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta, and Felice Antonio Merra. Leveraging content-style item representation for visual recommendation. In *ECIR (2)*, volume 13186 of *Lecture Notes in Computer Science*, pages 84–92. Springer, 2022.
- [24] Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. Assessing the impact of a user-item collaborative attack on class of users. In *Proceedings of the 1st Workshop on the Impact of Recommender Systems co-located with 13th ACM Conference on Recommender Systems, ImpactRS@RecSys 2019*, Copenhagen, Denmark, September 19, 2019, volume 2462 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2019.
- [25] Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. Adversarial machine learning in recommender systems (aml-recsys). In *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3-7, 2020*, pages 869–872. ACM, 2020.
- [26] Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. A survey on adversarial recommender systems: from attack/defense strategies to generative adversarial networks. *ACM Computing Surveys*, March 2021.
- [27] Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, and Felice Antonio Merra. How dataset characteristics affect the robustness of collaborative recommendation models. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 951–960. ACM, 2020.

- [28] Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, and Felice Antonio Merra. A regression framework to interpret the robustness of recommender systems against shilling attacks (discussion paper). In *IIR*, CEUR Workshop Proceedings. CEUR-WS.org, 2021.
- [29] Jacek Golebiowski, Felice Antonio Merra, Ziawasch Abedjan, and Felix Biessmann. Search filter ranking with language-aware label embeddings. In *WWW (Companion Volume)*, pages 121–125. ACM, 2022.
- [30] Felice Antonio Merra, Vito Walter Anelli, Tommaso Di Noia, Daniele Malitesta, and Alberto Mancino. Denoise to protect: a method to robustify visual recommenders from adversaries. In *Proceedings of the 46th International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2023, Taipei, July 2023*. ACM, 2023.
- [31] Felice Antonio Merra, Omar Zaidan, and Fabricio de Sousa Nascimento. Improving the relevance of product search for queries with negations. In *WWW (Companion Volume)*. ACM, 2023.
- [32] Tommaso Di Noia, Daniele Malitesta, and Felice Antonio Merra. Taamr: Targeted adversarial attack against multimedia recommender systems. In *50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops, DSN Workshops 2020, Valencia, Spain, June 29 - July 2, 2020*, pages 1–8. IEEE, 2020.

Berlin, April 11, 2023

Felice Antonio Merra