



Alessio Zanga

EDUCATION AND TRAINING

01/11/2021 – 17/02/2025

PH.D. IN COMPUTER SCIENCE F. Hoffmann - La Roche Ltd

Level in EQF EQF level 8

01/09/2019 – 30/07/2021 Milan, Italy

MASTER DEGREE IN COMPUTER SCIENCE (LM-18) University of Milano-Bicocca

Address Piazza dell'Ateneo Nuovo, 1, 20126, Milan, Italy | **Website** <https://en.unimib.it/> |

Field of study Information and Communication Technologies | **Final grade** 110/110 Laude | **Level in EQF** EQF level 7 |

Type of credits ECTS | **Number of credits** 120 | **Thesis** Causal Models for Microbial Communities in the Tropical Air Ecosystem

01/09/2016 – 30/07/2019 Milan, Italy

BACHELOR DEGREE IN COMPUTER SCIENCE (L-31) University of Milano-Bicocca

Address Piazza dell'Ateneo Nuovo, 1, 20126, Milan, Italy | **Website** <https://en.unimib.it/> |

Field of study Information and Communication Technologies | **Final grade** 110/110 Laude | **Level in EQF** EQF level 6 |

Type of credits ECTS | **Number of credits** 180 | **Thesis** Analysis and Classification of EEG Data

01/09/2011 – 30/07/2016 Caravaggio, Italy

HIGH SCHOOL LEAVING QUALIFICATION IN APPLIED SCIENTIFIC STUDIES Liceo Galileo Galilei

Address Via San Francesco n. 119/A, 24043, Caravaggio, Italy | **Website** <https://www.liceogalilei.edu.it/liceo-galilei-english> |

Field of study Natural sciences, mathematics and statistics | **Final grade** 100/100

WORK EXPERIENCE

□ **UNIVERSITY OF MILANO-BICOCCA – MILAN, ITALY**

ADJUNCT PROFESSOR – 01/10/2022 – 30/09/2023

Adjunct professor for the Master's Degree Course in Causal Networks (2223-2-F1801Q161). The course aims to provide a gentle introduction to causal inference and in particular to causal networks and structural causal models. In particular, the course gives strong motivations because, at the current state-of-the-art, modern machine learning experts need causality, and tools from causal modeling, to correctly address and effectively solve problems of decision making under uncertainty.

□ **UNIVERSITY OF MILANO-BICOCCA – MILAN, ITALY**

RESEARCH ASSISTANT – 01/08/2020 – 30/07/2021

Research assistant for developing a personalized ICT system using machine learning to support and monitor Mild Cognitive Impairment (MCI) patients. The SENIOR Project aims to develop and validate an ICT-based coaching system, rooted in Nudge theory, for supporting elderly individuals with MCI. The project involves a multi-center clinical trial with 200 seniors and uses machine learning to create personalized profiles, delivering tailored suggestions through an app. The expected outcomes include real-time monitoring to slow MCI progression, improved wellbeing for users, and the empowerment of caregivers through knowledge sharing, ultimately contributing to more sustainable healthcare



systems. Grant ID 2018-NAZ-0129 - "SENIOR - SystEm of Nudge theory based ICT applications for Older citizens" (ref. 2018-0826), Funder Entity: Fondazione Cariplo, CUP H45J18000590007 .

CONSULTING EXPERIENCE

14/10/2024 – 12/04/2025

Istituto Nazionale dei Tumori

Senior training consultant at the ABIGAIL (M) Project - Acquire BIGdata and Analytics for Innovation and Managerial Work - carried out in synergy with the management school Federmanager Academy in Rome and the Aldini Valeriani Foundation in Bologna, School of Industrial Management in Confindustria Emilia Area Centro. The latter will host workshops and experimental actions, making use of tutors identified by the partners involved.

01/06/2022 – 31/07/2022

Fondazione Aldini Valeriani

Research consultant with the aim of exploring the task of parameter learning in continuous networks under predefined equality and inequality constraint (i.e. any algebraic combination of such constraints). These functionalities rely on the bnlearn R package, and will also be available via a user-friendly Shiny interface.

Link <https://translate.google.com/translate?sl=it&tl=en&hl=en&u=https://www.fav.it/imprese/abigailm-acquisire-bigdata-e-analytics-per-linnovazione-e-il-lavoro-manageriale/&client=webapp>

01/11/2021 – 31/12/2021

Max Planck Institute for the Science of Human History

TEACHING EXPERIENCE

01/10/2024 – CURRENT

Master's Degree Course in Causal Networks (2425-2-F1801Q161)

The course aims to provide a gentle introduction to causal inference and in particular to causal networks and structural causal models. In particular, the course gives strong motivations because, at the current state-of-the-art, modern machine learning experts need causality, and tools from causal modeling, to correctly address and effectively solve problems of decision making under uncertainty.

Main contents are as follows: the potential outcome framework, main definitions and properties of probabilistic graphical models with specific reference to Bayesian networks, causal networks and structural causal models, randomized experiments, nonparametric identification of causal effect, estimation of causal effect, unobserved confounding, instrumental variables, structural learning from observational data and from observational and intervention data, basic concepts of transfer learning and transportability, and finally a basic introduction to counterfactuals.

01/10/2022 – 30/09/2023

Master's Degree Course in Causal Networks (2223-2-F1801Q161)

Link <https://elearning.unimib.it/course/info.php?id=44578#obiettivien>

VISITING PERIODS

01/03/2024 – 29/06/2024

Queen Mary University of London

Endometrial Cancer (EC) is cancer of the mucous lining, or endometrium, of the uterus. It is a common gynaecological disease affecting hundreds of thousands of women worldwide. Although most patients with EC are diagnosed at an early stage of the disease and have a favourable prognosis, approximately 90,000 patients around the world die every year because of EC. Surgery to remove the uterus (hysterectomy), possibly together with the ovaries (ovariectomy), is the typical initial treatment for EC. The choice of neo-adjuvant (pre-surgery) or adjuvant (post-surgery) treatments depends on patient outcome prognosis. The presence of pelvic and/or para-aortic lymph node metastases (LNM) is one of the most important prognostic factors for poor outcome. Therefore, the focus of this work is to develop a prognostic model capable of assessing the pre-operative risk of LNM in EC patients. Specifically, we want to learn a Causal Network (CN) from multi-centric data sets in presence of missing data and prior knowledge elicited from



experts' clinicians. This would enable us to model the cause-effect pairs in an explicit fashion, allowing clinicians to evaluate potential treatments plans according to underlying CN structure.

01/09/2023 – 01/12/2023

GIGA - Université de Liège

Inflammatory Bowel Disease (IBD) is a chronic inflammatory disorder of the gastro-intestinal tract. In this setting, the gut microbiome plays a crucial role in determining patients' response to a given therapy. To analyse the composition of the microbial community, DNA extraction and profiling was performed in different cohorts of IBD patients. In this work, we investigate the interplay between microbiota taxa and medical treatments using graphical models to discover the unknown interaction mechanisms, allowing us to explore both inter-individual variations and inter-taxa relationships. We leverage the framework of causal discovery to identify the effects of different therapies mediated by profiled microbiota taxa, incorporating experts' prior knowledge whenever available. Furthermore, it is possible to assess the impact of specific therapies with respect to patients' covariates, enabling experts to carefully evaluate personalized treatments on an individual level.

PUBLICATIONS

2025

Using GPT-4 to guide causal machine learning

Since its introduction to the public, ChatGPT has had an unprecedented impact. While some experts praised AI advancements and highlighted their potential risks, others have been critical about the accuracy and usefulness of Large Language Models (LLMs). In this paper, we are interested in the ability of LLMs to identify causal relationships. We focus on the well-established GPT-4 (Turbo) and evaluate its performance under the most restrictive conditions, by isolating its ability to infer causal relationships based solely on the variable labels without being given any other context by humans, demonstrating the minimum level of effectiveness one can expect when it is provided with label-only information. We show that questionnaire participants judge the GPT-4 graphs as the most accurate in the evaluated categories, closely followed by knowledge graphs constructed by domain experts, with causal Machine Learning (ML) far behind. We use these results to highlight the important limitation of causal ML, which often produces causal graphs that violate common sense, affecting trust in them.

2023

Causal Discovery with Missing Data in a Multicentric Clinical Study

Write here the description...Causal inference for testing clinical hypotheses from observational data presents many difficulties because the underlying data-generating model and the associated causal graph are not usually available. Furthermore, observational data may contain missing values, which impact the recovery of the causal graph by causal discovery algorithms: a crucial issue often ignored in clinical studies. In this work, we use data from a multi-centric study on endometrial cancer to analyze the impact of different missingness mechanisms on the recovered causal graph. This is achieved by extending state-of-the-art causal discovery algorithms to exploit expert knowledge without sacrificing theoretical soundness. We validate the recovered graph with expert physicians, showing that our approach finds clinically-relevant solutions. Finally, we discuss the goodness of fit of our graph and its consistency from a clinical decision-making perspective using graphical separation to validate causal pathways.

2023

Towards a Causal Decision-Making Framework for Recommender Systems

Causality is gaining more and more attention in the machine learning community and consequently also in recommender systems research. The limitations of learning offline from observed data are widely recognized, however, applying debiasing strategies like Inverse Propensity Weighting does not always solve the problem of making wrong estimates. This concept paper contributes a summary of debiasing strategies in recommender systems and the design of several toy examples demonstrating the limits of these commonly applied approaches. Therefore, we propose to map the causality frameworks of potential outcomes and structural causal models onto the recommender systems domain in order to foster future research and development. For instance, applying causal discovery strategies on offline data to learn the causal graph in order to compute counterfactuals or improve debiasing strategies.

2023

Towards a Transportable Causal Network Model Based on Observational Healthcare Data



Over the last decades, many prognostic models based on artificial intelligence techniques have been used to provide detailed predictions in healthcare. Unfortunately, the real-world observational data used to train and validate these models are almost always affected by biases that can strongly impact the outcomes validity: two examples are values missing not-at-random and selection bias. Addressing them is a key element in achieving transportability and in studying the causal relationships that are critical in clinical decision making, going beyond simpler statistical approaches based on probabilistic association. In this context, we propose a novel approach that combines selection diagrams, missingness graphs, causal discovery and prior knowledge into a single graphical model to estimate the cardiovascular risk of adolescent and young females who survived breast cancer. We learn this model from data comprising two different cohorts of patients. The resulting causal network model is validated by expert clinicians in terms of risk assessment, accuracy and explainability, and provides a prognostic model that outperforms competing machine learning methods.

2022

A Survey on Causal Discovery: Theory and Practice

Understanding the laws that govern a phenomenon is the core of scientific progress. This is especially true when the goal is to model the interplay between different aspects in a causal fashion. Indeed, causal inference itself is specifically designed to quantify the underlying relationships that connect a cause to its effect. Causal discovery is a branch of the broader field of causality in which causal graphs are recovered from data (whenever possible), enabling the identification and estimation of causal effects. In this paper, we explore recent advancements in causal discovery in a unified manner, provide a consistent overview of existing algorithms developed under different settings, report useful tools and data, present real-world applications to understand why and how these methods can be fruitfully exploited.

2022

Risk Assessment of Lymph Node Metastases in Endometrial Cancer Patients: A Causal Approach

Assessing the pre-operative risk of lymph node metastases in endometrial cancer patients is a complex and challenging task. In principle, machine learning and deep learning models are flexible and expressive enough to capture the dynamics of clinical risk assessment. However, in this setting we are limited to observational data with quality issues, missing values, small sample size and high dimensionality: we cannot reliably learn such models from limited observational data with these sources of bias. Instead, we choose to learn a causal Bayesian network to mitigate the issues above and to leverage the prior knowledge on endometrial cancer available from clinicians and physicians. We introduce a causal discovery algorithm for causal Bayesian networks based on bootstrap resampling, as opposed to the single imputation used in related works. Moreover, we include a context variable to evaluate whether selection bias results in learning spurious associations. Finally, we discuss the strengths and limitations of our findings in light of the presence of missing data that may be missing-not-at-random, which is common in real-world clinical settings.

2021

ACTA: A Mobile-Health Solution for Integrated Nudge-Neurofeedback Training for Senior Citizens

As the worldwide population gets increasingly aged, in-home telemedicine and mobile-health solutions represent promising services to promote active and independent aging and to contribute to a paradigm shift towards patient-centric healthcare. In this work, we present ACTA (Advanced Cognitive Training for Aging), a prototype mobile-health solution to provide advanced cognitive training for senior citizens with mild cognitive impairments. We disclose here the conceptualization of ACTA as the integration of two promising rehabilitation strategies: the "Nudge theory", from the cognitive domain, and the neurofeedback, from the neuroscience domain. Moreover, in ACTA we exploit the most advanced machine learning techniques to deliver customized and fully adaptive support to the elderly, while training in an ecological environment. ACTA represents the next-step beyond SENIOR, an earlier mobile-health project for cognitive training based on Nudge theory, currently ongoing in Lombardy Region. Beyond SENIOR, ACTA represents a highly-usable, accessible, low-cost, new-generation mobile-health solution to promote independent aging and effective motor-cognitive training support, while empowering the elderly in their own aging.

2020

An Attention-based Architecture for EEG Classification

Emerging studies in the deep learning community focus on techniques aimed to identify which part of a graph can be suitable for making better decisions and best contributes to an accurate inference. These researches (i.e., "attentional mechanisms" for graphs) can be applied effectively in all those situations in which it is not trivial to capture dependency between the involved entities while discharging useless information. This is the case, e.g., of functional connectivity in human brain, where rapid physiological changes, artifacts and high inter-subject variability usually



require highly trained clinical expertise. In order to evaluate the effectiveness of the attentional mechanism in such critical situation, we consider the task of normal vs abnormal EEG classification using brain network representation of the corresponding EEG recorded signals.

2020

Comparison of Attention-based Deep Learning Models for EEG Classification

Objective: To evaluate the impact on Electroencephalography (EEG) classification of different kinds of attention mechanisms in Deep Learning (DL) models. **Methods:** We compared three attention-enhanced DL models, the brand-new InstaGATs, an LSTM with attention and a CNN with attention. We used these models to classify normal and abnormal (ie, artifactual or pathological) EEG patterns. **Results:** We achieved the state of the art in all classification problems, regardless the large variability of the datasets and the simple architecture of the attention-enhanced models. We could also prove that, depending on how the attention mechanism is applied and where the attention layer is located in the model, we can alternatively leverage the information contained in the time, frequency or space domain of the dataset. **Conclusions:** with this work, we shed light over the role of different attention mechanisms in the classification of normal and abnormal EEG patterns. Moreover, we discussed how they can exploit the intrinsic relationships in the temporal, frequency and spatial domains of our brain activity. **Significance:** Attention represents a promising strategy to evaluate the quality of the EEG information, and its relevance, in different real-world scenarios. Moreover, it can make it easier to parallelize the computation and, thus, to speed up the analysis of big electrophysiological (eg, EEG) datasets.

CONFERENCES AND SEMINARS

10/09/2024 – 13/09/2024 Nijmegen, Netherlands

12th International Conference on Probabilistic Graphical Models (PGM 2024)

The International Conference on Probabilistic Graphical Models (PGM) is a biennial meeting that brings together researchers interested in all aspects of graphical models for probabilistic reasoning, decision making, and learning.

Link <https://www.ru.nl/en/about-us/events/probabilistic-graphical-models-pgm-2024>

09/11/2023 – 09/11/2023 Rome, Italy

2nd AIxIA Workshop on Artificial Intelligence For Healthcare (HC@AIxIA 2023)

The HC@AIxIA workshop aims at gathering researchers from academia, industry and medical centers for presenting and discussing the latest research results and ongoing works related to the application and impact of AI in the healthcare domain, to the larger extent, thus aiming at covering a wide spectrum of topics, including theoretical and practical aspects, methodologies, technologies, and systems.

Link <https://sites.google.com/unical.it/hcaixia2023>

06/11/2023 – 09/11/2023 Rome, Italy

22nd International Conference of the Italian Association for Artificial Intelligence (AIxIA 2023)

The 22nd International Conference of the Italian Association for Artificial Intelligence (AIxIA 2023) is organized by AIxIA (Associazione Italiana per l'Intelligenza Artificiale), which is a non-profit scientific society founded in 1988 and devoted to the promotion of Artificial Intelligence. The society aims to increase the public awareness of AI, encourage the teaching of it and promote research in the field.

Link <https://www.aixia2023.cnr.it/>

26/06/2023 – 27/06/2023 Lugano, Switzerland

1st Workshop on Sparse Inference on Complex Networks

Contributed talk on the published work "Causal Discovery with Missing Data in a Multicentric Clinical Study". This workshop aims to explore recent advances in sparse inference on complex networks. The study of network structures has become increasingly important in fields such as economics, epidemiology, ecology, and biology. However, analyzing complex networks remains a challenge, even with big data. This workshop will focus on three themes: causal regularization, diversification processes and random graphical models. Attendees can expect to hear keynote talks



from leading experts and engage in discussions on statistical inference for networks, graphical models, and machine learning methods.

Link <https://www.ci.inf.usi.ch/workshop-sparse-inference-on-complex-networks/>

12/06/2023 – 15/06/2023 Portoroz, Slovenia

21st International Conference of Artificial Intelligence in Medicine (AIME 2023)

The AIME 2023 conference will be hosted in person by University of Maribor Faculty of Health Sciences in Portorož, Slovenia from June 12-15, 2023. As previous AIME conferences (from Pavia in 1985 to Halifax in 2022), it will be a unique opportunity for researchers and other healthcare and AI experts to present significant theoretical, methodological and applied results related to the application of Artificial Intelligence (AI) in medicine.

Link <https://aime23.aimedice.info/>

28/11/2022 – 02/12/2022 Udine, Italy

21st International Conference of the Italian Association for Artificial Intelligence (AIxIA 2022)

The 21st International Conference of the Italian Association for Artificial Intelligence (AIxIA 2022) is organized by AIxIA (Associazione Italiana per l'Intelligenza Artificiale), which is a non-profit scientific society founded in 1988 and devoted to the promotion of Artificial Intelligence. The society aims to increase the public awareness of AI, encourage the teaching of it and promote research in the field.

Link <https://aixia2022.uniud.it/>

30/11/2022 – 30/11/2022 Udine, Italy

1st AIxIA Workshop on Artificial Intelligence For Healthcare (HC@AIxIA 2022)

The HC@AIxIA workshop aims at gathering researchers from academia, industry and medical centers for presenting and discussing the latest research results and ongoing works related to the application and impact of AI in the healthcare domain, to the larger extent, thus aiming at covering a wide spectrum of topics, including theoretical and practical aspects, methodologies, technologies, and systems.

Link <https://sites.google.com/unical.it/hcaixia2022>

21/08/2022 – 27/08/2022 Certosa di Pontignano, Castelnuovo Berardenga (Siena) Tuscany, Italy

5th Advanced Course on Data Science & Machine Learning (ACDL 2022)

The 5th Advanced Course on Data Science & Machine Learning (ACDL) is a full-immersion five-day Course at the Certosa di Pontignano (Castelnuovo Berardenga – Siena – Tuscany, Italy) on cutting-edge advances in Data Science and Machine Learning with lectures delivered by world-renowned experts. The Course provides a stimulating environment for academics, early career researches, Post-Docs, PhD students and industry leaders. Participants will also have the chance to present their results with talks or posters, and to interact with their colleagues, in a convivial and productive environment.

Link <https://acd12022.icas.cc/>

23/02/2020 – 25/02/2020 La Valletta, Malta

13th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSIGNALS 2020)

The purpose of the International Conference on Bio-inspired Systems and Signal Processing is to bring together researchers and practitioners from multiple areas of expertise, including biology, medicine, engineering and as well as computer and data science, interested in studying and using models and techniques inspired from or applied to biological systems. A diversity of signal types can be found in this area, including video, audio, electrophysiological signals, medical imaging, and other biological sources of information. The analysis and use of these signals is a multidisciplinary area including signal processing, pattern recognition and artificial intelligence techniques, amongst others.

Link <http://www.biosignals.biostec.org/?y=2020>



NETWORKS AND MEMBERSHIPS

01/01/2022 – CURRENT Washington, DC, USA
International Biometric Society (IBS)

Devoted to the Development and Application of Statistical and Mathematical Theory and Methods in the Biosciences.

01/01/2020 – CURRENT Istituto di Scienze e Tecnologie della Cognizione, Via San Martino della Battaglia n.44, Rome
The Italian Association for Artificial Intelligence (AIxIA)

Link <https://aixia.it/en/>

01/09/2019 – CURRENT Dipartimento di Informatica Sistemistica e Comunicazione, Viale Sarca, n. 336, Milan
Models and Algorithms for Data & Text Mining (MADLab)

A Research Laboratory at DISCo of the University of Milano-Bicocca on Data and Text Mining, Probabilistic Graphical Models, Bayesian Networks and Continuous Time Bayesian Networks.

Link <https://mad.disco.unimib.it/>

LANGUAGE SKILLS

Mother tongue(s): ITALIAN

Other language(s):

	UNDERSTANDING		SPEAKING		WRITING
	Listening	Reading	Spoken production	Spoken interaction	
ENGLISH	C1	C1	C1	C1	C1

Levels: A1 and A2: Basic user; B1 and B2: Independent user; C1 and C2: Proficient user

DRIVING LICENCE

Driving Licence: B

| 03/02/2016

SKILLS

Microsoft Office | Microsoft Word | Microsoft Excel | Microsoft Powerpoint | Outlook | Google Docs | Google Drive | LinkedIn | Skype | Zoom | Motivated | Good listener and communicator | Team-work oriented | Organizational and planning skills | Decision-making | Analytical skills | Critical thinking | Creativity | Reliability