

Curriculum Vitae et Studiorum

<i>Name</i>	Andrea Borghesi
<i>Date of Birth</i>	January 24 1988
<i>Citizenship</i>	Italian
<i>Affiliation</i>	DISI, University of Bologna
<i>Address</i>	Viale del Risorgimento, 2, 40136 Bologna (BO) - Italy
<i>Phone</i>	+39 051 20 93274
<i>E-mail</i>	andrea.borghesi3@unibo.it
<i>Web Page</i>	https://www.unibo.it/sitoweb/andrea.borghesi3
<i>ORCID ID</i>	0000-0002-2298-2944

Short Bio

Andrea Borghesi received the M.S. degree in Computer Engineering and the Ph.D. in Computer Engineering from the University of Bologna (Italy), in 2013 and 2017 respectively. He is currently an *Assistant Professor, tenure-track (“RTD-B”)* at the Department of Computer Science and Engineering (DISI) of the University of Bologna. He received the *Italian National Scientific Qualification for the role of Associate Professor in Computer Engineering* (ASN 2018-2020 settore 09/H1 II Fascia, VI Quadrimestre). He held the role of Executive Scientific Representative (HPC scientific unit) at the Inter-departments Center for Artificial Intelligence at University of Bologna, Italy (ALMA-AI, in Italian: “Centro Interdipartimentale di Intelligenza Artificiale”).

His research interests are focused on **optimization techniques and Machine Learning approaches for complex systems and novel computer paradigms**, with a special attention towards power and energy related aspects and predictive and prescriptive maintenance. His expertise spans different areas, with the general target of minimizing power consumption and automatize and improve the maintenance of the target systems: optimization and implementation of novel scheduling policies for supercomputers; creation of Machine/Deep Learning predictive models for automated maintenance; design and deployment of fine-grained monitoring infrastructure for heterogeneous large computing machines; studying the interaction of methods from different Artificial Intelligence areas to solve problems over complex, real-world systems. A big focus of his research has been the exploration of areas at the intersection of different branches of Artificial Intelligence, such as combining Machine Learning, optimization and simulation models.

Bologna, 17 February 2026:

Contents

Position & Education	2
Professional Activities	3
Teaching Activities	14
Academic Service Activities	17
Research Interests	18
Publication List	21

Position & Education

Record of Employment

- 11/2052 – *Present*
Research Consultant and Project Coordinator (focus on AI compliance) at the DIGIS Center Fondazione Bruno Kessler, Trento, Italy, University of Bologna, Italy
- 10/2025 – *Present*
Associate Professor at the Department of Information Engineering and Computer Science (DISI), University of Bologna, Italy
- 03/2020 – *Present*
Executive Scientific Representative (HPC scientific unit) at the Inter-departments Center for Artificial Intelligence at University of Bologna, Italy (ALMA-AI, in Italian: “Centro Interdipartimentale di Intelligenza Artificiale”)
- 12/2022 – 10/2025
Assistant Professor (RTDB) at the Department of Information Engineering and Computer Science (DISI), University of Bologna, Italy
- 2019 – *2022*
Assistant Professor (RTDA) at the Department of Information Engineering and Computer Science (DISI), University of Bologna, Italy
- 2018 – *2019*
Postdoctoral fellow at the Department of Information Engineering and Computer Science (DISI), University of Bologna, Italy
Supervisor: Prof. Michela Milano
- 2017 – 2018
Research intern at the Department of Electrical, Electronic and Information Engineering (DEI),

University of Bologna, Italy
Supervisor: Prof. Luca Benini

- 02/2016 – 07/2016
Research Fellow at University College of Cork (UCC), Ireland
Supervisor: Prof. Barry O’Sullivan
- June 2013 – December 2013
Research intern at the Department of Information Engineering and Computer Science (DISI), University of Bologna, Italy
Supervisor: Prof. Michela Milano

Education

- **Ph.D. in Computer Science and Engineering**, University of Bologna, 2017
Thesis Title: Power-aware Job Dispatching in High Performance Computing Systems
Advisors: Prof. Michela Milano, Prof. Luca Benini
 - **Master Degree in Computer Engineering**, University of Bologna, 2013
Score: 110/110 cum laude
Thesis Title: Integrazione di ottimizzazione e simulazioni per il piano energetico regionale dell’Emilia-Romagna (Eng: Integrating optimization and simulation approaches to devise the energy plan of the Emilia-Romagna Region)
Advisor: Prof. Michela Milano
Topic: Applications of Artificial Intelligence
-

Professional Activities

Contribution to National & International Academic Research Projects

INTERNATIONAL PROJECTS:

- **HORIZON-EUROHPC-JU-2023-ENERGY-04-HORIZON-JU-RIA-101177590-SEANERGYYS: Software for Efficient and Energy-Aware Supercomputers**
Period: 05/2025 – 04/2029
Role: Principal Investigator (Involvement: HIGH)
Budget: 319571.88 euros
Description:
SEANERGYYS creates an integrated European software solution that optimises resource utilisation and reduces the energy used for real-world workload mixes. It, therefore, improves the throughput of HPC systems, generating more R&D results for a given energy budget. The solution comprises a comprehensive monitoring infrastructure, an Artificial Intelligence data analytics system, and a dynamic scheduling and resource management system.
EUROPEAN-URBAN-INITIATIVE-TALEA: Green cells leading the Green transition
Period: 12/2024 – 11/2028

Role: Principal Investigator (Involvement: HIGH)

Budget: 26149.00 euros

Description:

Bologna and Emilia-Romagna face rising temperatures, making climate resilience a priority. TALEA introduces modular units integrating nature-based solutions, technology, and citizen science to address Urban Heat Islands and Heat Waves. These units are adaptable and replicable, deployed in participatory-selected areas to provide climate shelters and co-designed services, fostering social, cultural, and environmental benefits. Data and results will integrate into Bologna's Digital Twin.

- **DIGITAL-2021-CLOUD-AI-01-101083674-Pre-PAI: PREPARATION FOR THE DEVELOPMENT OF THE AI-ON-DEMAND PLATFORM**

Period: 11/2022 – 08/2023

Role: Team member (Involvement: LOW/MEDIUM)

Description:

Pre-PAI will be the realization of “Preparatory actions for the AI-on-demand platform” providing the blueprint for the further development, deployment, and operation of the European AI-on-demand platform, realized in “Deployment of the AI-on-demand platform”. The project will carry out a comprehensive requirement analysis for different stakeholder groups, mainly SMEs, industrial sectors, and public administration. These requirements analyses will lead to an overall roadmap and plan to build and consolidate the AI-on-demand platform. It includes a comprehensive concept to provide trustworthy AI assets into the platform to support European leadership on trustworthy AI and to offer AI “Made in Europe”

- **HORIZON-CL4-2021-HUMAN-01-02-101070000-AI4EUROPE: AN AI ON-DEMAND PLATFORM TO SUPPORT RESEARCH EXCELLENCE IN EUROPE**

Period: 12/2022 – 11/2023

Role: Team member (Involvement: MEDIUM/HIGH)

Description:

The main objective of AI4Europe is to support and facilitate a sustainable digital platform and experimentation environment through the creation of open research channels and mechanisms that foster the European AI research ecosystem, academic and industrial, and that maximises the academic, social, and industrial impact while it seamlessly integrates other projects, platforms, and solutions.

- **HORIZON-CL4-2021-HUMAN-01-24-101070363-AEQUITAS: ASSESSMENT AND ENGINEERING OF EQUITABLE, UNBIASED, IMPARTIAL AND TRUSTWORTHY AI SYSTEMS**

Period: 11/2022 – 10/2025

Role: Risk Manager/WP Leader (Involvement: HIGH)

Description:

AI-based decision support systems are increasingly deployed in industry, in the public and private sectors, and in policy-making. As our society is facing a dramatic increase in inequalities and intersectional discrimination, we need to prevent AI systems from amplifying this phenomenon and rather mitigate it. To trust these systems, domain experts and stakeholders need to trust the decisions. How the principles underpinning fairness and trustworthiness translate into technical, social, and lawful requirements in the AI system design is still an open question. Similarly, we don't know how to test if a system is compliant with these principles and repair it.

AEQUITAS proposes the design of a controlled experimentation environment for developers and users to create controlled experiments for 1) assessing the bias in AI systems, e.g., identifying potential causes of bias in data, algorithms, and interpretation of results, 2) providing effective methods and engineering guidelines to repair, remove, and mitigate bias, 3) provide fairness-by-design guidelines, and software engineering techniques to design new bias-free systems.

- **H2020-ICT-2020-101017142-StairwAI: Stairway to AI: Ease the Engagement of Low-Tech users to the AI-on-Demand platform through AI**

Period: 01/2021 – 12/2023

Role: Technical Contributor (Involvement: HIGH)

Website: <https://cordis.europa.eu/project/id/101017142>

Description:

The StairwAI project targets low-tech users to facilitate their engagement on the AI4EU on-demand Platform. This will be achieved through a new service layer, enriching the functionalities of the on-demand platform and containing: (1) a multi-lingual interaction layer enabling conversations with the Platform in the user's own language, (2) a horizontal matchmaking service for the automatic discovery of AI assets (tools, data sets, AI experts, consultants, papers, courses, etc.) meeting the user business needs and (3) a vertical matchmaking service that will dimension and provision hardware resources through a proper hardware provider (HPC, Cloud and Edge infrastructures).

Andrea Borghesi is directly involved in different areas of the project, with a special focus on the work packages 5 and 6, relative to the horizontal and vertical matchmaking; in particular, he is leading T6.3. In this area, his contributions revolve around the combination of Machine Learning and optimization to guide the matching process, e.g. by integrating data-driven models in constraint programming models for hardware dimensioning and algorithm configuration. He also contributed to writing the project proposal.

- **H2020-ICT-2019-952215-TAILOR: Foundations of Trustworthy AI - Integrating Reasoning, Learning and Optimization**

Period: 09/2020 – 07/2024

Role: Technical Contributor (Involvement: MEDIUM)

Website: <https://tailor-network.eu/>

Description:

Maximising opportunities and minimising risks associated with artificial intelligence (AI) requires a focus on human-centred trustworthy AI; this can be achieved by collaborations between research excellence centres. Currently, this work is carried out by an isolated scientific community where research groups are working individually or in smaller networks; TAILOR aims to bring these groups together in a single scientific network on the Foundations of Trustworthy AI.

Andrea Borghesi research concerns the integration of two AI paradigms, namely Machine Learning and optimization, for instance by embedding empirical models (derived from data) into combinatorial optimization models. In the context of TAILOR, the research activity concerns the development of methods to integrate these paradigms with a special focus towards trustworthy AI (especially Task 4.2).

- **H2020-ICT-2018-3-857191-IoTwins: Distributed Digital Twins for industrial SMEs: a big-data platform**

Period: 09/2019 – 08/2022

Role: WP Leader (Involvement: HIGH)

Website: <https://www.iotwins.eu/>

Description:

The concept of digital twins has been around, but the Internet of Things has enabled its cost-effective implementation. Digital twins refer to a virtual representation of a physical product or process. The EU-funded IoTwins project plans to build testbeds for digital twins in the manufacturing and facility management sectors. The digital models will integrate data from various sources such as data APIs, historical data, embedded sensors and open data.

Andrea Borghesi is involved in IoTwins as leader of the work package 3. His research activity focuses on Machine Learning services for digital twins (both in large data centers and the manufacturing area), with the task of collecting requirements, developing and integrating ML methods and approaches, deploying them and evaluating their implementation – within work package 3. He is also involved in the development of a vertical, holistic, monitoring infrastructure to collect and store heterogeneous data from large data centers (big data storing and processing), in addition to analysing the collected data and developing predictive models for optimized maintenance.

- **H2020-FETPROACT-732631-OPRECOMP: Open transprecision computing**

Period: 02/2019 – 12/2019

Role: Technical Contributor (Involvement: MEDIUM/HIGH)

Website: <http://oprecomp.eu/>

Description:

This project explores a novel computing paradigm (transprecision computing) which aims at improving the energy efficiency of computing systems by reducing the precision of single variables and intermediate operations to provide the minimum accuracy required for the final results (in contrast with the “conservative” precision provided by standard type systems).

The research activity of Andrea Borghesi is focused on approaches for the fine-tuning of the precision of Floating Point variables, exploiting techniques from the artificial intelligence area, from machine learning to linear/integer optimization (mathematical/constraint programming). He collaborates especially in work packages 6.1 and 6.2 and is part of the UNIBO representatives at official review meetings.

- **FP7-RI-312763-H2020-653838-PRACE-3IP: OpenPower cluster**

Period: 01/2017 – 12/2018

Role: Technical Contributor (Involvement: HIGH)

Website: <http://www.prace-ri.eu/prace-3ip/>

Description:

PRACE (Partnership for Advanced Computing in Europe) started the process for a pre-commercial procurement on Whole-System Design for Energy Efficient High Performance Computing (HPC) in 2012 with a workshop on the usage of Pre-Commercial Procurement (PCP) in HPC, to obtain R&D services which should result in future PRACE HPC systems becoming more energy efficient, especially exploiting OpenPower technologies.

Andrea Borghesi participated in the design of a novel High Performance Computing system, the D.A.V.I.D.E. supercomputer, in collaboration with industrial and academic partners. His activity was mainly focused on: 1) developing real-time job dispatching policies to limit the system power consumption under a given threshold (*power capping*); 2) building an extensive data collection infrastructure to gather performance counters and job information on a heterogeneous supercomputer with a very high granularity and with big data capabilities; 3) characterizing the HPC system

behaviour through machine learning models built with the data collected from the monitoring infrastructure, in particular with the purpose of anomaly detection and power consumption prediction.

- **FP7-IDEAS-ERC-291125-MULTITHERMAN: Multi-Scale Thermal Management of Computing Systems**

Period: 01/2014 – 03/2018

Role: Technical Contributor (Involvement: HIGH)

Website: <http://projects.eees.dei.unibo.it/multitherman/>

Description:

This project studies novel strategies for thermal planning/management which are alternatives to worst-case design; the main aim is to pave the way towards methodologies enabling lower power/energy consumption and thermal optimization. The development of a synergistic performance, power and thermal management strategy requires major breakthroughs in several areas, namely architectures, run-time systems, resource management middleware, code optimization tools and programming models.

The major contributions of Andrea Borghesi in the context of the project fall in three areas: 1) the development of novel scheduling policies aimed at curtailing power consumption in large-scale systems and at improving system performance (w.r.t. both users and owners/administrators satisfaction), ranging from the algorithm design to the actual implementation/deployment of the proposed solution; 2) the creation of thermal/power models of complex systems (HPC machines and data centers) to be used for predictive tasks (mainly machine learning models trained in supervised fashion); 3) the creation of data collection frameworks to gather information from heterogeneous sources (for instance, in the context of HPC machines, measures coming from the computing nodes and related components).

- **FP7-ICT-2011-7-288147-ePolicy: Engineering the POLicy-making Life CYcle**

Period: 01/2013 – 09/2014

Role: Technical Contributor (Involvement: LOW/MEDIUM)

Website: <http://www.epolicy-project.eu/>

Description:

The project's main aim is to support policy makers in their decision process across a multi-disciplinary effort aimed at engineering the policy making life-cycle. For the first time, global and individual perspectives on the decision processes are merged and integrated. The project focuses on regional planning and promotes the assessment of economic, social and environmental impacts during the policy-making process (at both the global and individual levels)

Andrea Borghesi followed two main research avenues: 1) understanding the impact of public economic incentives on the decision process of private citizens using agent-based models, with an explicit focus on the adoption of photovoltaic panels in the Italian Emilia-Romagna region; 2) devising novel incentive schemes to foster the fulfillment of macro-level objectives acting only on micro-level stakeholders (through optimization techniques).

NATIONAL PROJECTS:

- **CINECA-ISCRA-C-HP10CDLFLZ-SKIDD-AI: Investigating Fairness In Large-Lange Models**

Period: 11/2025 – 08/2026

Role: Principal Investigator (Involvement: HIGH)

Budget: 80000 compute core hours on HPC Tier-0 systems

Website: <http://www.hpc.cineca.it/services/iscra/>

Description:

AI-based skin disease diagnosis shows potential for improving healthcare access, but current fairness solutions are still exploratory and perform poorly on real-world, non-dermoscopic images. Our work proposes a two-stage pipeline—data augmentation using diffusion models for better skin tone representation and deep learning classification—explicitly addressing fairness and performance on imbalanced, low-quality datasets.

- **CINECA-ISCRA-C-HP10CDLFLZ-FairLLM: Investigating Fairness In Large-Lange Models**

Period: 08/2024 – 07/2025

Role: Principal Investigator (Involvement: HIGH)

Budget: 82000 compute core hours on HPC Tier-0 systems

Website: <http://www.hpc.cineca.it/services/iscra/>

Description:

The project is composed of the following tasks: 1. a benchmark of several state-of-the-art LLMs in terms of fairness and their broader performance; 2. a taxonomy of the different types of bias in state-of-the-art LLMs; 3. an analysis of the trade-off between fairness and performance; 4. a set of tools and approaches to mitigate the bias in LLMs, to obtain fair models.

- **CINECA-ISCRA-C-HP10CFXHZU-Pro-job: Predictive Profiling and Models of HPC Workload**

Period: 06/2023 – 05/2024

Role: Principal Investigator (Involvement: HIGH)

Budget: 24000 compute core hours on HPC Tier-0 systems

Website: <http://www.hpc.cineca.it/services/iscra/>

Description:

The project aims to extract knowledge from large amounts of data collected from production supercomputers (in particular, M100) and to build predictive models for the workload typically encountered in such supercomputers. Such models will focus on predicting job failures and studying the relationship between job execution and HPC nodes anomalies. The overarching purpose of the project is the creation of tools which could be integrated in a digital twin of the HPC system to improve its reliability and facilitate its maintenance.

- **CINECA-ISCRA-C-HP10CVV2VY-vertOpt: Vertical Match Making of Artificial Intelligence Applications**

Period: 09/2021 – 06/2022

Role: Principal Investigator (Involvement: HIGH)

Budget: 16000 compute core hours on HPC Tier-0 systems

Website: <http://www.hpc.cineca.it/services/iscra/>

Description:

This project main objective is the creation of a data set which will be used to develop an optimization engine for vertical matchmaking of Artificial Intelligence algorithms. Given a set of requirements from the end users and a problem to be solved, vertical matchmaking consists in finding the appropriate set of hardware (HW) resources and algorithm (together with its optimal configuration) for the target task.

- **CINECA-ISCRA-C-HP10CVDID8-V-Match: Vertical Match Making of Artificial Intelligence Applications**

Period: 07/2020 – 04/2021

Role: Principal Investigator (Involvement: HIGH)

Budget: 17406 compute core hours on HPC Tier-0 systems

Website: <http://www.hpc.cineca.it/services/iscra/>

Description:

The project's main goal is the design and development of a vertical matchmaking service for Artificial Intelligence applications. Given a set of requirements from the end users - time constraints, costs, privacy, etc. – vertical matchmaking consists of finding the appropriate set of hardware (HW) resources resources that satisfy a given algorithm and tool. The aim of the proposed activity is to perform a preliminary exploration of the vertical matchmaking.

- **CINECA-ISCRA-C-HP10C0V1Z5-EXADATA: Exascale Monitoring and Big Data Analytics**

Period: 10/2019 – 07/2021

Role: Principal Investigator (Involvement: HIGH)

Budget: 37500 compute core hours on HPC Tier-0 systems

Website: <http://www.hpc.cineca.it/services/iscra/>

Description:

The project's main aim is the extraction and analysis of the wealth of data generated by a monitoring infrastructure in a supercomputer, with the goal of improved system performance and maintenance and towards data center automation. The target HPC system is the Marconi supercomputer, hosted by the CINECA consortium and monitored via a holistic infrastructure called Examon. The objectives of the research activity were threefold: 1) develop of the digital twin of a supercomputer in production; 2) devise a data-driven, approach for anomaly detection in an HPC system; 3) create power/thermal/energy/performance predictive models.

RESEARCH PROJECTS IN COLLABORATION WITH PRIVATE INSTITUTIONS AND COMPANIES :

- **Job-scheduling optimization for High-Performance Computing systems, commissioned research by ENI**

Period: Set 2022 - March 2023

Role: Principal Investigator (Involvement: HIGH)

Budget: 20000 euros

Description:

The research consists of collaborating with ENI system administrators and another consulting company (DataReply) for the development of optimized algorithms for job scheduling and allocation in High-Performance Computing (HPC) systems. ENI has a great interest in this domain, as they own the largest HPC systems in Italy. The activity focuses on creating the optimization model for the HPC job dispatching solution (based on Quadratic Unconstrained Binary Optimization) and assisting in the creation of a prototype to be adopted in ENI's machines.

- **Advanced diagnostic in complex industrial systems, with YANMAR Italy S.P.A. *Period:* Jun 2019 - Feb 2020**

Role: Scientific Responsible (Involvement: HIGH)

Description:

The research consisted of the development of a diagnostic framework for power generation and

electric/thermal (both hot and cold) energy consumption devices, monitored through suitable sensors. These devices can undergo unexpected changes in their behaviour. The goals of the activity concerned 1) the definition of general (e.g. non device-specific) Machine Learning (ML) models to support the diagnosis of anomalies and faulty states and 2) the experimental evaluation of such models.

INTERNATIONAL COLLABORATIONS:

- **Energy Efficient High Performance Computing Group (EEHPCWG)**

Period: 06/2016 – present

Role: Member *Website:* <https://eehpcwg.llnl.gov/>

Description:

The work group's mission is to encourage the implementation of energy conservation measures, energy efficient design in high performance computing (HPC), and share ideas.

Andrea Borghesi participated to the group meetings and discussions, collaborating in the writing of white papers, international conference papers, and technical reports.

Program Chair & Organization Committees

Andrea Borghesi has been **chair/co-chair** of:

- Organizing Chair of the 2nd International Workshop on Technologies for AI Governance (TAIG 2025), Oct 2025 2025, (https://www.ln.edu.hk/sds/news-and-events-v0/sds_news_and_event_2025_TAIG) co-located with ECAI 2025, Bologna, Italy
- Organizing Chair of the 1st International Workshop on the Environmental Sustainability of High-Performance Software (SHiPS23), Nov 2023, (<https://site.unibo.it/ships23/en>) co-located with Supercomputing 2023 (Denver, US)
- Local Chair of Computing Frontiers 2023, the 20th ACM International Conference on Computing Frontiers 8-11 May 2023, Bologna, Italy
- Co-chair of HiPEAC 2020, the International Conference on High Performance and Embedded Architectures and Compilers, Jan 2020, Bologna, Italy

Andrea Borghesi has been **scientific panel member** of:

- IS CRA initiative – IS CRA is the Italian SuperComputing Resource Allocation promoted by CINECA, the Italian most powerful HPC center, which directly awards in excess of 400 millions core processor hours, to ensure an adequate supply to scientists and engineers for HPC-related research.
- LISA initiative – LISA is the Interdisciplinary Laboratory for Advanced Simulation and it aims at the promotion of research projects with an high original value with, if possible, a return on the Lombard territory and an industrial interest.

Andrea Borghesi has served as **guest editor** for:

- Applied Sciences Journal, ISSN 2076-3417, (MDPI), special issue "Machine Learning for Industry 4.0: From Manufacturing and Embedded Systems to Cloud Computing and Data Centers", May 2020 - Feb 2021 (https://www.mdpi.com/journal/applsci/special_issues/Machine_Learning_Industry_4)

Andrea Borghesi has been a member of the **program committee** of:

- International Joint Conference on Artificial Intelligence (IJCAI), since 2017
- European Conference on Artificial Intelligence (ECAI), since 2019
- Association for the Advancement of Artificial Intelligence Conference (AAAI), since 2018
- International Conference on Principles and Practice of Constraint Programming (CP), since 2018
- ACM International Conference on Computing Frontiers (CF), since 2021
- European Congress Embedded Real Time Systems (ERTS), since 2019

Andrea Borghesi has been a member of the following **international research collaborations** of:

- Energy Efficient High Performance Computing Working Group, since 2017

Presentations at Workshops & Conferences

Andrea Borghesi has been **invited speaker** at the following workshops and conferences (national and international):

- Workshop Digital Twins for HPC, co-hosted with Supercomputing Conference 2024 (SC2024), title of the intervention "Building a Digital-Twins from Scratch - The Case Study of CINECA Datacenter", 17 Nov 2024, Atlanta (GA, US)
- AI for manufacturing Workshop organized within the Cluster-Mech General Assembly (2024), title of the intervention "AI e Manifattura: sinergie per l'Efficienza e l'Innovazione", Jun 27 2024, Bologna (Italy)
- Digital twins for Industry 4.0, workshop organized by the Confindustria Emilia-Romagna association, title of the intervention "Industria 4.0 e gemelli digitali: applicazioni di Big Data e Intelligenza Artificiale per creare modelli predittivi e ottimizzazione", September 27 2023, Forlì (Italy)
- Borderless Body step #2: dancing human and artificial creativity, workshop organized by DAMSLab, April 20, 2022, Bologna (Italy), <https://site.unibo.it/damslab/it/eventi/borderless-body-step-2-dancing-1>
- Openday IoTwins, workshop organized by CINECA, title of the intervention "Examion e la manutenzione predittiva in ambito HPC", Jun 28 2022, Bologna (Italy)
- Data Week 2022, Big Data Value Association (BDVA), title of the intervention "IoTwins - Big Data Platform for optimized and replicable Industrial and Facility Management models ", May 26 2022, Online
- Workshop "Evento Gemelli Digitali - Applicazioni Pratiche e Case History" (English translation: "Digital twins Event - Practical Applications and Case History"), May 25, 2022, Bologna (Italy)
- CINECA School "Digital twins per l'industria" (English translation: "Digital twins for the industry"), Dec 1,4 2021, Bologna (Italy)
- European Big Data Value Forum (EBDV), title of the intervention "IoTwins - Distributed Digital Twins for industrial SMEs: a big-data platform", Nov 29 2021, Ljubljana (Slovenia) & Online

- Workshop “Intelligenza Artificiale per il settore Manifatturiero” (English translation: “Artificial Intelligence for the manufacturing sector”) hosted at MECSPE-2019, Mar 30, 2019, Parma (Italy)
- CINECA Workshop “HPC for Industry 4.0”, May 21-23 2019, Milan (Italy)
- PRACE School “Energy Efficiency in HPC”, organized at CINECA, Bologna (Italy), Feb 22, 2019
- Workshop “Presentation of D.A.V.I.D.E. supercomputer”, organized at CINECA, Bologna (Italy), Nov 29, 2018
- Workshop “Deep Learning and Machine Learning for HPC”, organized at CINECA, Bologna (Italy), Oct 23, 2018

Andrea Borghesi **chaired panels** at the following workshops and conferences (national and international):

- Chair of the panel "Panel - Fairness in Large Language Models: A Linguistic Perspective", part of the workshop "AI Fairness Event: Addressing Bias, Stereotypes, and Discrimination in AI Systems", Bologna (Italy), Dec 06 2024

Andrea Borghesi **presented his research** at the following workshops and conferences (national and international):

- International Workshop on Artificial Intelligence and Creativity, co-located with ECAI 2024 (CREAI24), Oct 20, 2024, Santiago De Compostela (Spain)
- International Workshop on Health Intelligence (W3PHIAI 2024), co-located with AAAI 24, Feb 1, 2024, Vancouver (Canada)
- The 35th IEEE International Conference on Tools with Artificial Intelligence (ICTAI), Nov 7, 2023, Atlanta (US)
- 32nd International Joint Conference on Artificial Intelligence (IJCAI-23), Aug 24, 2023, Macao (SAR, China)
- XXII Conference of the Italian Association for Artificial Intelligence (AI*IA2022), Nov 29, 2022, Udine (Italy)
- ACM International Conference on Computing Frontiers (CF20), May 11, 2020, Catania (Italy) – The conference took place remotely due to Covid19 pandemic emergency
- XVIII Conference of the Italian Association for Artificial Intelligence (AI*IA2019), Nov 21, 2019, Rende (Italy)
- CINECA workshop “Energy Efficiency in HPC”, Feb 22, 2019, Bologna (Italy)
- 1st IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS2019), Mar 19, 2019, Hsinchu (Taiwan)
- Thirsty-First Annual Conference on Innovative Applications of Artificial Intelligence (IAAI-19), Jan 29, 2019, Honolulu (US)
- CINECA Workshop “Deep Learning and Machine Learning for HPC”, Oct 23, 2018, Bologna (Italy)

- ACM International Conference on Computing Frontiers (CF18), May 8, 2018, Ischia (Italy)
- ISC High Performance (ISC16), Jun 20, 2016, Frankfurt (Germany)
- XIV Conference of the Italian Association for Artificial Intelligence (AI*IA2015), Sep 24, 2015, Ferrara (Italy)
- 21st International Conference on Principles and Practices of Constraint Programming (CP15), Sep 3, 2015, Cork (Ireland)
- International Conference on High Performance Computing and Simulation (HPCS15), Jul 21, 2015, Amsterdam (Netherlands)
- International Conference on Renewable Energy Research and Application (ICRERA2014), Oct 20, 2014, Milwaukee (US)
- 20th International Conference on Principles and Practices of Constraint Programming (CP14), Sep 10, 2014, Lyon (France)
- XIII Conference of the Italian Association for Artificial Intelligence (AI*IA2013), Dec 4, 2013, Turin (Italy)
- 19th International Conference on Principles and Practices of Constraint Programming (CP13), Sep 19, 2013, Uppsala (Sweden)
- 27th European Conference on Modeling and Simulation 2013 (ECMS13), May 29, 2013, Alesund (Norway)

Referee Services in Journals & Conferences

Andrea Borghesi a **reviewer** for the following journals (J) and PC/SPC/APC for the following conferences (C):

- (J) ACM Computing Surveys
- (J) Journal of Artificial Intelligence Research (JAIR)
- (J) IEEE Transactions on Dependable and Secure Computing (TDSC)
- (J) IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)
- (J) IEEE Transactions on Parallel and Distributed Systems (TPDS)
- (J) Elsevier Engineering Applications of Artificial Intelligence
- (J) Elsevier Environmental Science & Policy
- (J) Advanced Engineering Informatics
- (J) Springer Wireless Networks
- (J) Elsevier ICT Express

- (J) Springer Constraints
- (J) Intelligenza Artificiale
- (C) IJCAI International Joint Conference on Artificial Intelligence
- (C) AAAI Conference on Artificial Intelligence
- (C) ECAI European Conference on Artificial Intelligence
- (C) Hawaii International Conference on System Sciences (HICSS)
- (C) International Conference on Principles and Practice of Constraint Programming (CP)
- (C) International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)
- (C) International Conference on Computing Frontiers (CF)
- (C) European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)
- (C) IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing (CCGrid)
- (C) European Congress Embedded Real Time Systems (ERTS)

Prizes and Awards

Andrea Borghesi has received the following prizes & awards:

- Winner of the best paper award at the W3PHIAI Workshop, held in conjunction with AAAI-24. Paper titled “Generation of Clinical Skin Images with Pathology with Scarce Data” (February 2024)
- Winner of the prize "IJCAI 2023 – Distinguished Program Committee Member"
- Winner of the prize "IJCAI 2019 – Distinguished Program Committee Member"
- Winner of GAUSS AWARD 2016 (Research Paper Award), obtained with the paper “Predictive Modeling for Job Power Consumption in HPC Systems”

Teaching Activities

In the academia (bachelor & master degree)

2021 – *present*

Adjunct Professor

Master’s degree in Engineering Management, University of Bologna, Italy

Metodi e Modelli di Data Analytics, 3 CFUs (30 hours)

Lectures, exercises, final project supervision

Academic years: 2021-2022 (Spring), 2022-2023 (Fall), 2023-2024 (Fall)

Total number of hours (in the classroom or lab): 90

2020 – *present*

Adjunct Professor

Bachelor's degree in Computer Engineering, University of Bologna, Italy

Fondamenti di Informatica (Modulo 2), 6 CFUs (60 hours)

Lectures, exercises, final project supervision

Academic years: 2020-2021 (Fall), 2021-2022 (Fall), 2022-2023 (Fall), 2023-2024 (Fall)

Total number of hours (in the classroom or lab): 240

2017 – 2020

Adjunct Professor

Bachelor's degree in Automation/Electrical Engineering, University of Bologna, Italy

Fondamenti di Informatica (Modulo 2), 3CFUs (30 hours)

Lectures, exercises, final project supervision

Academic years: 2017-2018 (Spring), 2018-2019 (Spring), 2019-2020 (Spring)

Total number of hours (in the classroom or lab): 90

In the academia (PhD courses)

2020 – *present*

Course title: From embedded devices to the cloud: a vertical perspective on IoT systems (10 hours)

PhD in Data science and computation, University of Bologna, Italy

Courses held: June 2020, July 2022, July 2023

Total number of hours (in the classroom or lab): 30

2021 – *present*

Course title: From embedded devices to the cloud: a vertical perspective on IoT systems (10 hours)

PhD in Computer Science and Engineering, University of Bologna, Italy

Courses held: July 2021

Total number of hours (in the classroom or lab): 10

In the academia (Teaching Assistant/Tutor/Mentor)

2014 – *present*

Teaching Assistant

Master degree Computer/Electronics Engineering, University of Bologna, Italy

Intelligent Systems

Seminars, exercises, final project supervision

2013 – 2017

Teaching Assistant

Bachelor degree Automation/Electrical Engineering, University of Bologna, Italy

Fondamenti di Informatica

Lectures, exercises, final project supervision

2019 – *present*

Teaching Assistant

Master degree Electronics Engineering, University of Bologna, Italy

Lab of Big Data Architectures M

Lectures, projects, final project supervision

2017 – 2018

Teaching Assistant

Master degree Advanced Automotive Engineering, University of Bologna, Italy

Algorithms and Systems for Big Data Processing

Lectures

Outside of the Academia - Other Courses

- Course held at OPTIT SRL (Cesena, Italy), titled “Advanced techniques and applications of Deep Learning”, 20 hours, April 2021
- Course held at HPE-COXA (Modena Italy), titled “ML and AI for predictive analytics”, 16 hours, 8 and 16 October 2020
- Course held at OPTIT SRL (Cesena, Italy), titled “Introduction to Machine and Deep Learning”, 4 hours, 29 November 2019

Student Supervision Activity

Andrea Borghesi has supervised or co-supervised the following **Ph. D. Students**:

- Davide Leone; PhD Program in Data Science and Computation; 36 cycle, University of Bologna
- Chiara Bellatreccia; National PhD Program in AI; 41 cycle, University of Bologna and Fondazione Bruno Kessler
- Ali Tariq Nagi; National PhD Program in AI; 40 cycle, University of Bologna and University of Pisa
- Emmen Farooq; PhD Program in Computer Science and Engineering; 37 cycle, University of Bologna
- Francesco Antici; PhD Program in Computer Science and Engineering; 37 cycle, University of Bologna
- Martin Molan; PhD Program in Data Science and Computation; 36 cycle, University of Bologna

Andrea Borghesi has supervised the following 6 **post-graduate students and temporary research fellows**:

- Emmen Farooq, temporary research fellow, Jun 2025 - Dec 2025, research title: “Federated Learning and Transfer Learning Approaches for Predictive Maintenance in High-Performance Computing Systems”, University of Bologna
- Andrea Virgillito, temporary research fellow, Feb 2025 - Jan 2026, research area: “Development and Deployment of AI-based Models for Predictive Maintenance of Rotating Equipment”, University of Bologna

- Carmine Di Santi, temporary research fellow, May 2022 - September 2023, research area: “Development and Deployment of AI-based Models for Industrial Digital Twins”, University of Bologna
- Matteo Francobaldi, temporary research fellow, Jan 2022 - Oct 2023, research area: “Vertical optimization methods for Hardware dimensioning and AI algorithms configuration”, University of Bologna
- Andrea Boscarino, temporary research fellow, Nov 2020 -Dec 2021, research area: “Learning models and optimization algorithms for static variable precision tuning in the transprecision computing area”, University of Bologna
- Marco Di Felice, post-graduate research scholarship, Jun 2019 - Feb 2020, research area: “Advanced diagnostics in complex industrial systems”, in collaboration with YANMAR Italy, University of Bologna

Andrea Borghesi has supervised dozens of students during their textbfinal graduation project (19 bachelor’s students and 49 master’s students), as well as several students during their internships (currently 11). He also supervised 2 students during their “**Summer of Code HPC**” **internship**, performed in conjunction between the University of Bologna and the CINECA consortium. Finally, he has supervised 6 students during their “**Erasmus+**” **internship**, hosted by the University of Bologna.

Academic Service Activities

I have performed the following service activities in academia (not including the student supervision during final projects and thesis):

2025 – *present*

Member of the Inter-department Committee of Technological Area ("Commissione interdipartimentale di A
University of Bologna, Italy

2025

PhD Selection Committee Member for the "Dottorato Nazionale", 41° Cycle

2023 – *present*

Executive Scientific Representative (HPC scientific unit)

Inter-departments Center for Artificial Intelligence at the University of Bologna, Italy

2025 – *present*

Member of the **Teaching/Curriculum Committee** ("Commissione Didattica")

Bachelor’s and Master’s degrees in Management Engineering, University of Bologna, Italy

2023 – *present*

Member of the **Teaching/Curriculum Committee** ("Commissione Didattica")

Bachelor’s and Master’s degrees in Computer Engineering, University of Bologna, Italy

2024 – *present*

Member of the **Teaching/Curriculum Committee** ("Commissione Didattica")
Master's degree in Computer Engineering, University of Bologna, Italy

Research Interests

My main research interests cover different areas of the field of Artificial Intelligence (AI), prevalently Machine Learning (ML) and optimization and simulation with Agent-Based Models (ABM), with a focus on applied solutions and especially in the areas of management of large computing systems and power/energy aware topics. Broadly speaking, I have always been attracted to demonstrate the potential of AI applications in many fields [J16].

A core research direction I have explored consists of exploring the areas at the intersection of multiple branches of AI research, with a particular interest in mingling the sub-symbolic layers with higher-level, human-like, reasoning-based approaches[C8]. For instance, I have tackled the challenge of integrating data-driven models for predictive and classification tasks, such as ML models, into optimization and simulation models, striving to combine the benefits of the two paradigms. In an orthogonal direction, I have also studied the impact of using constraints derived from the optimization area to boost the performance of ML models (e.g., [C35]).

Hardware dimensioning The problem of determining what is the right hardware (HW on premise or the cloud) architecture and its dimensioning for AI algorithms is still crucial. Searching for the optimal solution is often challenging, as it is not trivial to anticipate the behaviour of an algorithm on diverse architectures. I tackled this problem by proposing an automated decision support tool called HADA and described in [J11, C23]. I have exploited the Empirical Model Learning paradigm, based on the integration of ML models into an optimization problem. The key idea is to integrate domain knowledge held by experts with data-driven models that learn the relationships between HW requirements and AI algorithm performances. In particular, the approach starts with benchmarking multiple AI algorithms on different HW resources, generating data used to train ML models; then, optimization is used to find the best HW configuration that respects user-defined constraints (e.g., budget, time, solution quality)[C27].

Prescriptive and predictive maintenance in industrial contexts and large-scale data centers

Another research direction stemming from my interest in complex systems (such as High Performance Computing, HPC, systems, data centers, factories, etc.) explores the possibility of improving and automating the maintenance of these systems, both from a predictive (e.g. forecast when a fault is going to happen) and a prescriptive (decide a set of actions to prevent/mitigate the fault) point of view, for example, as I have studied in the context of Industry 4.0[C16]. The key in this area is building accurate models capable of representing the behaviour of complex systems.

The first step to create useful models is to have good data describing the system in its entirety. For this purpose, I have participated in the design and deployment of a supercomputer with an innovative, scalable, and fine-grained monitoring infrastructure, capable of recreating an extremely accurate image of the monitored system[C1, C6, C4]. As the culmination of several years of collaboration with CINECA, the Italian supercomputing center, I have collected fine-grained data detailing the behaviour of the entire HPC facility; this data has then been made public to foster the advancement of research in the HPC domain

[J7] (this is the largest public data set describing the behaviour of a real, production supercomputer over multiple years). Similarly, other data collection campaigns I have collaborated with focused on specific widely-used HPC applications, to characterize their energy profile [J2].

Thanks to the data collected from holistic monitoring systems, I have developed an approach for building models describing the state of the system. In large systems, it is not easy to identify malfunctions in real time; as part of my research, I have developed automated tools for anomaly detection, based on ML/DL models and using HPC machines as a case study. I experimented mainly with two broad families of ML models: I) supervised ones[C44, J19], where the ML models need to be trained with a balanced data set containing examples belonging to all status-classes to detect (thus anomalies need to be injected into the system to obtain the training set[C43]); II) semi-supervised models, where the ML model learns the normal behaviour of a supercomputer and can then detect anomalies since they diverge from the normal distribution [C11, J4] – the last tool is also publicly available as a prototype¹.

As a follow-up to the semi-supervised approach, I also deployed an anomaly detection model on a real supercomputer to perform online anomaly detection in real time, with minimal overhead[C17, C3]. Additionally, I have also dealt with the problem of *explaining* the output of fault detection models in the HPC domain[C39, C40]. Along the same line, I have also contributed to the development of a fully unsupervised approach for anomaly detection[J18], with the great benefit of not requiring labels at training time (annotated data is quite scarce in the HPC domain). Additionally, I have also explored the potential of Federated Learning (FL) approaches to improve anomaly detection in HPC[C25, J13, C26, J14], as well as the transferability of these approaches[J15].

Moving beyond pure anomaly detection, the natural follow-up is to study the possibility of *anticipating* the insurgence of faults, a piece of information which could be used in predictive maintenance solutions. In this direction, I proposed a DL-based approach combining a semi-supervised model and a supervised one to anticipate the anomalies in a production tier-0 supercomputer [J9, C41, C42]. In the direction of anomaly prediction, I’ve also recently contributed to a line of research employing graphical neural networks to better model the behaviour of racks of nodes in a supercomputer[C38, J17].

I have also collaborated in the creation of a suite of AI-based services for optimizing the management of the supercomputer facility, in a series of tasks ranging from anomaly detection in computing nodes and failure predictions of hard drives, to power and thermal models of the computing nodes, and the server rooms; this work culminated in a tool called *ExaMon-X*[J6]. Again, in the context of fault prediction but in a different industrial domain, namely, failure prediction of ball bearings in an industrial machine producing bottle caps, I have devised and validated a DL model for remaining useful life estimation[C32].

Transprecision computing Transprecision computing is a novel computing paradigm aimed at designing systems to deliver just the required precision for intermediate computations rather than tolerating errors implied by imprecise HW or SW computations. Adopting this principle, the user-defined constraints on the precision of final results are always met, but intermediate operations can be deployed to custom, lower precision compute units to save energy. Within the scope of the *OPRECOMP* project, I contributed with an exploration of approaches for mixed precision tuning of Floating Point (FP) variables, i.e. fine-tune the precision by specifying the number of bits to be used for the mantissa and the exponent of the FP variable (using a smaller number of bits decreases the precision, potentially saving energy)[C21].

In this context, I developed a prototype tool for static precision tuning, *StaticFPTuner*², which combines ML models (to represent and learn the relation between variable precision and overall accuracy of the computation) and optimization via Mathematical Programming (to find the optimal precision given

¹https://github.com/AndreaBorghesi/anomaly_detection_HPC

²<https://github.com/oprecomp/StaticFPTuner>

a target accuracy). In the context of transprecision computing, I also studied the benefits attainable by injecting prior domain knowledge into DL models, obtaining a significant increase in prediction accuracy[C9].

Optimization and power-aware job dispatchers for High Performance Computing systems

High Performance Computing (HPC) systems are large machines composed of thousands of heterogeneous components. Improving their performance and facilitating their maintenance has been a very interesting research challenge over the years. For instance, determining the optimal scheduling order for submitting jobs is a complex task that I have tackled using Constraint Programming and ad-hoc heuristic algorithms cite bartolini2014proactive, galleguillos2017data. If we consider also the power aspect and we aim for a power-aware job dispatcher, the difficulty of the problem increases: first, we need to develop predictive models capable of estimating a job’s power consumption before its execution[C10, C2, J1], then AI-inspired approaches can be combined with heuristic techniques to devise optimal scheduling policies[C15, C14, J3] while respecting the tight time limits arising from the real-time requirements of job dispatching in HPC systems

I also studied the optimal computation of the frequency assignment for a set of jobs already running on a supercomputer in the event of a sudden reduction in available power cite borghesi2019freqAssign. Dealing with a related aspect, I’ve also been involved in the development of a runtime framework for parallel applications capable of dynamically reducing the frequency of computing nodes in HPC systems during communication phases, which are not impacted by the frequency reduction[J10]. An additional related issue I addressed is the design of pricing schemes to foster the adoption of power-aware strategies, schemes capable of benefiting both system owners and users[J5]. Strictly connected to energy considerations and optimization of job dispatchers more broadly, the estimation of the duration of HPC applications is crucial, as more informed decisions can lead to much improved scheduling policies; for this reason, I worked on estimating the duration of HPC jobs using DL approaches[J12, C36].

Remaining in the area of energy efficiency for HPC and data center, over the years I have participated in the Energy Efficient High Performance Computing Working Group (EEHPCWG³), whose effort I have helped with surveys and detailed analysis, documenting the state of the art and state of practice in both academia and industry in terms of energy efficient solutions[C37, C34].

Agent-Based Models for energy policies Designing and evaluating energy policies is a difficult challenge because the energy sector is a complex system that cannot be adequately understood without using models that merge economic, social, and individual perspectives. Policy makers need precise models to estimate the impact of their plans. In my research, I focused on policies for promoting energy production from renewable energy sources and, in particular, on photovoltaic (PV) power generation in the Emilia-Romagna region[C20, C18]. In this context, I worked on ABMs developed by merging self-reported behavior (interviews with PV installers) and observed behavior (the historical trend of installed PV power). In this way, the agent decision process is grounded on realistic assumptions, and the historical data can be used to fine-tune the models, obtaining a method capable of predicting the PV power adoption [J8, C33].

Artificial Intelligence & Arts AI has demonstrated to be very relevant in the creative domain as well, being used to generate music, images, text, videos, and other artistic products traditionally considered an exclusive domain of the human mind. In this context, AI should not necessarily be seen as an adversary, but rather its usage can be integrated into the human artistic creation to explore novel possibilities

³<https://eehpcwg.llnl.gov/>

and generate useful insights about the human creative process itself. To this end, I have worked on AI approaches generating synthetic music combining constraint programming and deep learning [C30, C46]. Beyond the music domain, AI is very useful in the context of dance choreographies. In this field, I proposed a methodological approach for the artistic evaluation of both AI and human artistic creations in the field of robotic dance [C31, C24].

AI & Fairness AI has the potential to reproduce bias and discrimination present in nowadays society. I have worked to cope with this problem, to detect and mitigate bias in AI algorithms (e.g., in the educational [C29] or health sector [C13]). My main contribution in this field is the design of a Fair-by-Design methodology to help developers in creating fair applications, in conjunction with social and legal experts [C22]

In a similar direction, I worked on the generation of synthetic images of skin patches with diseases, starting with very few training samples. This is relevant to fairness as current training sets of skin disease images tend to under-represent minorities – and thus it is very hard to train a fair ML classifier on them [C12]. This is part of a larger research avenue covering the intersection of AI and health systems [J20, C7].

Publication List

Journals

- [J1] Francesco Antici, Andrea Borghesi, Zeynep Kiziltan, Jens Domke, and Andrea Bartolini. An online algorithm for power consumption prediction of hpc workload. *Future Generation Computer Systems*, page 108064, 2025.
- [J2] Marcello Artioli, Andrea Borghesi, Marta Chinnici, Anna Ciampolini, Michele Colonna, Davide De Chiara, and Daniela Loreti. C6enpls: A high-performance computing job dataset for the analysis of linear solvers’ power consumption. *Future Internet*, 17(5):203, 2025.
- [J3] Andrea Borghesi, Andrea Bartolini, Michele Lombardi, Michela Milano, and Luca Benini. Scheduling-based power capping in high performance computing systems. *Sustainable Computing: Informatics and Systems*, 19:1–13, 2018.
- [J4] Andrea Borghesi, Andrea Bartolini, Michele Lombardi, Michela Milano, and Luca Benini. A semisupervised autoencoder-based approach for anomaly detection in high performance computing systems. *Engineering Applications of Artificial Intelligence*, 85:634–644, 2019.
- [J5] Andrea Borghesi, Andrea Bartolini, Michela Milano, and Luca Benini. Pricing schemes for energy-efficient hpc systems: Design and exploration. *The International Journal of High Performance Computing Applications*, 33(4):716–734, 2019.
- [J6] Andrea Borghesi, Alessio Burrello, and Andrea Bartolini. Examon-x: a predictive maintenance framework for automatic monitoring in industrial iot systems. *IEEE Internet of Things Journal*, 2021.

- [J7] Andrea Borghesi, Carmine Di Santi, Martin Molan, Mohsen Seyedkazemi Ardebili, Alessio Mauri, Massimiliano Guarrasi, Daniela Galetti, Mirko Cestari, Francesco Barchi, Luca Benini, et al. M100 exadata: a data collection campaign on the cineca’s marconi100 tier-0 supercomputer. *Scientific Data*, 10(1):288, 2023.
- [J8] Andrea Borghesi and Michela Milano. Merging observed and self-reported behaviour in agent-based simulation: A case study on photovoltaic adoption. *Applied Sciences*, 9(10):2098, 2019.
- [J9] Andrea Borghesi, Martin Molan, Michela Milano, and Andrea Bartolini. Anomaly detection and anticipation in high performance computing systems. *IEEE Transactions on Parallel and Distributed Systems*, 2021.
- [J10] Daniele Cesarini, Andrea Bartolini, Andrea Borghesi, Carlo Cavazzoni, Mathieu Luisier, and Luca Benini. Countdown slack: a run-time library to reduce energy footprint in large-scale mpi applications. *IEEE Transactions on Parallel and Distributed Systems*, 31(11):2696–2709, 2020.
- [J11] Allegra De Filippo, Andrea Borghesi, Andrea Boscarino, and Michela Milano. Hada: An automated tool for hardware dimensioning of ai applications. *Knowledge-Based Systems*, page 109199, 2022.
- [J12] Allegra De Filippo, Emanuele Di Giacomo, and Andrea Borghesi. Machine learning approaches to predict the execution time of the meteorological simulation software cosmo. *Journal of Intelligent Information Systems*, pages 1–25, 2024.
- [J13] Emmen Farooq, Michela Milano, and Andrea Borghesi. Harnessing federated learning for anomaly detection in supercomputer nodes. *Future Generation Computer Systems*, 2024.
- [J14] Emmen Farooq, Michela Milano, and Andrea Borghesi. Federated lstm autoencoders for time series anomaly detection in production-scale hpc systems. *Knowledge-Based Systems*, page 115043, 2025.
- [J15] Emmen Farooq, Michela Milano, and Andrea Borghesi. Federated transfer learning for anomaly detection in hpc systems: First real-world validation on a tier-0 supercomputer. *Expert Systems with Applications*, page 129754, 2025.
- [J16] Gregor Molan, Gregor Dolinar, Jovan Bojkovski, Radu Prodan, Andrea Borghesi, and Martin Molan. Model for quantitative estimation of functionality influence on the final value of a software product. *IEEE Access*, 2023.
- [J17] Martin Molan, Mohsen Seyedkazemi Ardebili, Junaid Ahmed Khan, Francesco Beneventi, Daniele Cesarini, Andrea Borghesi, and Andrea Bartolini. Graafe: Graph anomaly anticipation framework for exascale hpc systems. *Future Generation Computer Systems*, 2024.
- [J18] Martin Molan, Andrea Borghesi, Daniele Cesarini, Luca Benini, and Andrea Bartolini. Ruad: unsupervised anomaly detection in hpc systems. *Future Generation Computer Systems*, 2022.
- [J19] Alessio Netti, Zeynep Kiziltan, Ozalp Babaoglu, Alina Sîrbu, Andrea Bartolini, and Andrea Borghesi. A machine learning approach to online fault classification in hpc systems. *Future Generation Computer Systems*, 2019.
- [J20] Daniele Zama, Andrea Borghesi, Alice Ranieri, Elisa Manieri, Luca Pierantoni, Laura Andreozzi, Arianna Dondi, Iria Neri, Marcello Lanari, and Roberta Calegari. Perspectives and challenges of telemedicine and artificial intelligence in pediatric dermatology. *Children*, 11(11):1401, 2024.

Conferences & Workshops

- [C1] Wissam Abu Ahmad, Andrea Bartolini, Francesco Beneventi, Luca Benini, Andrea Borghesi, Marco Cicala, Privato Forestieri, Cosimo Gianfreda, Daniele Gregori, Antonio Libri, et al. Design of an energy aware petaflops class high performance cluster based on power architecture. In *2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 964–973. IEEE, 2017.
- [C2] Francesco Antici, Andrea Borghesi, Jens Domke, and Zeynep Kiziltan. Uopc: A user-based online framework to predict job power consumption in hpc systems. In *ISC High Performance 2025 Research Paper Proceedings (40th International Conference)*, pages 1–12. Prometheus GmbH, 2025.
- [C3] Francesco Antici, Andrea Borghesi, and Zeynep Kiziltan. Online job failure prediction in an hpc system. In *European Conference on Parallel Processing*, pages 167–179. Springer, 2023.
- [C4] Andrea Bartolini, Francesco Beneventi, Andrea Borghesi, Daniele Cesarini, Antonio Libri, Luca Benini, and Carlo Cavazzoni. Paving the way toward energy-aware and automated datacentre. In *Proceedings of the 48th International Conference on Parallel Processing: Workshops*, page 8. ACM, 2019.
- [C5] Andrea Bartolini, Andrea Borghesi, Thomas Bridi, Michele Lombardi, and Michela Milano. Proactive workload dispatching on the eurora supercomputer. In *International Conference on Principles and Practice of Constraint Programming*, pages 765–780. Springer, 2014.
- [C6] Andrea Bartolini, Andrea Borghesi, Antonio Libri, Francesco Beneventi, Daniele Gregori, Simone Tinti, Cosimo Gianfreda, and Piero Altoè. The davide big-data-powered fine-grain power and performance monitoring support. In *Proceedings of the 15th ACM International Conference on Computing Frontiers*, pages 303–308. ACM, 2018.
- [C7] Chiara Bellatreccia, Daniele Zama, Arianna Dondi, Luca Pierantoni, Laura Andreozzi, Iria Neri, Marcello Lanari, Andrea Borghesi, and Roberta Calegari. Ai-based skin disease diagnosis with non-dermoscopic images: Tackling bias and data limitations. 2026.
- [C8] Alice Bizzarri, Andrea Borghesi, Samuele Bortolotti, Luc De Raedt, Devendra Dhama, Michelangelo Diligenti, Sebastijan Dumančić1, Boi Faltings11, Elisabetta Gentili, Alfonso Gerevini12, et al. Benchmarking in neuro-symbolic ai. In *Learning and Reasoning: 4th International Joint Conference on Learning and Reasoning, IJCLR 2024, and 33rd International Conference on Inductive Logic Programming, ILP 2024, Nanjing, China, September 20–22, 2024, Proceedings*, page 238. Springer Nature, 2025.
- [C9] Andrea Borghesi, Federico Baldo, Michele Lombardi, and Michela Milano. Injective domain knowledge in neural networks for transprecision computing. In *International Conference on Machine Learning, Optimization, and Data Science*, pages 587–600. Springer, 2020.
- [C10] Andrea Borghesi, Andrea Bartolini, Michele Lombardi, Michela Milano, and Luca Benini. Predictive modeling for job power consumption in hpc systems. In *International conference on high performance computing*, pages 181–199. Springer, 2016.
- [C11] Andrea Borghesi, Andrea Bartolini, Michele Lombardi, Michela Milano, and Luca Benini. Anomaly detection using autoencoders in high performance computing systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9428–9433, 2019.

- [C12] Andrea Borghesi and Roberta Calegari. Generation of clinical skin images with pathology with scarce data. In *AI for Health Equity and Fairness: Leveraging AI to Address Social Determinants of Health*, pages 47–64. Springer, 2024.
- [C13] Andrea Borghesi, Roberta Calegari, et al. Ai for health equity and fairness. *STUDIES IN COMPUTATIONAL INTELLIGENCE*, pages 47–64, 2024.
- [C14] Andrea Borghesi, Francesca Collina, Michele Lombardi, Michela Milano, and Luca Benini. Power capping in high performance computing systems. In *International Conference on Principles and Practice of Constraint Programming*, pages 524–540. Springer, 2015.
- [C15] Andrea Borghesi, Christian Conficoni, Michele Lombardi, and Andrea Bartolini. Ms3: A mediterranean-style job scheduler for supercomputers-do less when it’s too hot! In *2015 International Conference on High Performance Computing & Simulation (HPCS)*, pages 88–95. IEEE, 2015.
- [C16] Andrea Borghesi, Giuseppe Di Modica, Paolo Bellavista, Varun Gowtham, Alexander Willner, Daniel Nehls, Florian Kintzler, Stephan Cejka, Simone Rossi Tisbeni, Alessandro Costantini, et al. Iotwins: Design and implementation of a platform for the management of digital twins in industrial scenarios. In *2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, pages 625–633. IEEE, 2021.
- [C17] Andrea Borghesi, Antonio Libri, Luca Benini, and Andrea Bartolini. Online anomaly detection in hpc systems. In *Proceedings of the 2019 IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS 2019)*, pages 229–233. IEEE, 2019.
- [C18] Andrea Borghesi and Michela Milano. Multi-agent simulator of incentive influence on pv adoption. In *2014 International Conference on Renewable Energy Research and Application (ICRERA)*, pages 556–560. IEEE, 2014.
- [C19] Andrea Borghesi, Michela Milano, and Luca Benini. Frequency assignment in high performance computing systems. In *International Conference of the Italian Association for Artificial Intelligence*, pages 151–164. Springer, 2019.
- [C20] Andrea Borghesi, Michela Milano, Marco Gavanelli, and Tony Woods. Simulation of incentive mechanisms for renewable energy policies. In *ECMS*, pages 32–38, 2013.
- [C21] Andrea Borghesi, Giuseppe Tagliavini, Michele Lombardi, Luca Benini, and Michela Milano. Combining learning and optimization for transprecision computing. In *Proceedings of the 17th ACM International Conference on Computing Frontiers*, pages 10–18, 2020.
- [C22] G Ciatto, M Matteini, L Sartori, M Rebrean, C Muller, A Borghesi, R Calegari, et al. Ai-fairness: the fairbridge approach to practically bridge the gap between socio-legal and technical perspectives. In *PROCEEDINGS OF THE ANNUAL HAWAII INTERNATIONAL CONFERENCE ON SYSTEM SCIENCES*, pages 6504–6513. IEEE Computer Society, 2025.
- [C23] Allegra De Filippo and Andrea Borghesi. Constrained hardware dimensioning for ai algorithms. In *PAIS 2022*, pages 145–148. IOS Press, 2022.

- [C24] Allegra De Filippo, Luca Giuliani, Eleonora Mancini, Andrea Borghesi, Paola Mello, Michela Milano, et al. Towards symbiotic creativity: A methodological approach to compare human and ai robotic dance creations. In *Thirty-Second International Joint Conference on Artificial Intelligence*, pages 5806–5814, 2023.
- [C25] Emmen Farooq and Andrea Borghesi. A federated learning approach for anomaly detection in high performance computing. In *2023 IEEE 35th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 496–500. IEEE, 2023.
- [C26] Emmen Farooq and Andrea Borghesi. Lstm-based unsupervised anomaly detection in high-performance computing: A federated learning approach. In *2024 IEEE International Conference on Big Data (BigData)*, pages 7735–7744. IEEE, 2024.
- [C27] Matteo Francobaldi, Allegra De Filippo, Andrea Borghesi, Nikola Pizurica, Igor Jovančević, Tim Llewellynn, and Miguel de Prado. Tinderai: Support system for matching ai algorithms and embedded devices. In *The International FLAIRS Conference Proceedings*, volume 36, 2023.
- [C28] Cristian Galleguillos, Alina Sirbu, Zeynep Kiziltan, Ozalp Babaoglu, Andrea Borghesi, and Thomas Bridi. Data-driven job dispatching in hpc systems. In *International Workshop on Machine Learning, Optimization, and Big Data*, pages 449–461. Springer, 2017.
- [C29] Joseph Giovanelli, Matteo Magnini, Liam James, Giovanni Ciatto, Angel S Marrero, Andrea Borghesi, Gustavo A Marrero, and Roberta Calegari. Unfair inequality in education: A benchmark for ai-fairness research. *Unfair Inequality in Education: A Benchmark for AI-Fairness Research*, 2024.
- [C30] Luca Giuliani, Francesco Ballerini, Allegra De Filippo, and Andrea Borghesi. Musicomb: a sample-based approach to music generation through constraints. In *2023 IEEE 35th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 194–198. IEEE, 2023.
- [C31] Luca Giuliani, Allegra De Filippo, Andrea Borghesi, Paola Mello, Michela Milano, et al. A multi-modal perspective for the artistic evaluation of robotic dance performances. In *Proceedings of the First Workshop on Artificial Intelligence and Creativity*, pages 84–93, 2022.
- [C32] Bassem Hichri, Anass Driate, Andrea Borghesi, and Francesco Giovannini. Predictive maintenance based on machine learning model. In Ilias Maglogiannis, Lazaros Iliadis, John Macintyre, and Paulo Cortez, editors, *Artificial Intelligence Applications and Innovations*, pages 250–261, Cham, 2022. Springer International Publishing.
- [C33] Valerio Iachini, Andrea Borghesi, and Michela Milano. Agent based simulation of incentive mechanisms on photovoltaic adoption. In *Congress of the Italian Association for Artificial Intelligence*, pages 136–148. Springer, 2015.
- [C34] Gregory A Koenig, Matthias Maiterth, Siddhartha Jana, Natalie Bates, Kevin Pedretti, Milos Puzovic, Andrea Borghesi, Andrea Bartolini, and Dave Montoya. Energy and power aware job scheduling and resource management: Global survey—an in-depth analysis. In *The 2nd International Industry/University Workshop on Data-center Automation, Analytics, and Control (DAAC’18), in conjunction with Supercomputing’18*, pages 2018–11.

- [C35] M. Lombardi, F. Baldo, A. Borghesi, and M. Milano. An analysis of regularized approaches for constrained machine learning. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12641 LNAI:112–119, 2021.
- [C36] Daniela Loreti, Davide Leone, and Andrea Borghesi. Duration-informed workload scheduler. In *International Conference on High Performance Computing*, pages 3–14. Springer, 2025.
- [C37] Matthias Maiterth, Gregory Koenig, Kevin Pedretti, Siddhartha Jana, Natalie Bates, Andrea Borghesi, Dave Montoya, Andrea Bartolini, and Milos Puzovic. Energy and power aware job scheduling and resource management: Global survey—initial analysis. In *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 685–693. IEEE, 2018.
- [C38] Martin Molan, Junaid Ahmed Khan, Andrea Borghesi, and Andrea Bartolini. Graph neural networks for anomaly anticipation in hpc systems. In *Companion of the 2023 ACM/SPEC International Conference on Performance Engineering*, pages 239–244, 2023.
- [C39] Martin Molan, Andrea Borghesi, Francesco Beneventi, Massimiliano Guarrasi, and Andrea Bartolini. An explainable model for fault detection in hpc systems. In *International Conference on High Performance Computing*, pages 378–391. Springer, 2021.
- [C40] Martin Molan, Andrea Borghesi, Luca Benini, and Andrea Bartolini. Analysing supercomputer nodes behaviour with the latent representation of deep learning models. In *European Conference on Parallel Processing*, pages 171–185. Springer, 2022.
- [C41] Martin Molan, Andrea Borghesi, Luca Benini, and Andrea Bartolini. Semi-supervised anomaly detection on a tier-0 hpc system. In *Proceedings of the 19th ACM International Conference on Computing Frontiers*, CF '22, page 203–204, New York, NY, USA, 2022. Association for Computing Machinery.
- [C42] Martin Molan, Andrea Borghesi, Luca Benini, and Andrea Bartolini. Semi-supervised anomaly detection on a tier-0 hpc system. In *Proceedings of the 19th ACM International Conference on Computing Frontiers*, pages 203–204, 2022.
- [C43] Alessio Netti, Zeynep Kiziltan, Ozalp Babaoglu, Alina Sîrbu, Andrea Bartolini, and Andrea Borghesi. Finj: A fault injection tool for hpc systems. In *European Conference on Parallel Processing*, pages 800–812. Springer, 2018.
- [C44] Alessio Netti, Zeynep Kiziltan, Ozalp Babaoglu, Alina Sîrbu, Andrea Bartolini, and Andrea Borghesi. Online fault classification in hpc systems through machine learning. In *European Conference on Parallel Processing*, pages 3–16. Springer, 2019.
- [C45] Stefano Teso, Laurens Bliiek, Andrea Borghesi, Michele Lombardi, Neil Yorke-Smith, Tias Guns, and Andrea Passerini. Machine learning for combinatorial optimisation of partially-specified problems: Regret minimisation as a unifying lens. *arXiv preprint arXiv:2205.10157*, 2022.
- [C46] Lorenzo Tribuiani, Luca Giuliani, Allegra De Filippo, and Andrea Borghesi. Expert-musicomb: Injective domain knowledge in a neuro-symbolic approach for music generation. In Allegra De Filippo, François Pachet, Valentina Presutti, and Luc Steels, editors, *Proceedings of the 3rd Workshop on Artificial Intelligence and Creativity co-located with 27th European Conference on Artificial Intelligence (ECAI 2024)*, Santiago de Compostela, Spain, October 20, 2024, volume 3810 of *CEUR Workshop Proceedings*, pages 46–58. CEUR-WS.org, 2024.