Computer Science Dept.
University of Turin, Italy
iacopo.colonnelli@unito.it
glassofwhiskey.github.io
GlassOfWhiskey
0000-0001-9290-2017
yFutTsMAAAAJ

Iacopo Colonnelli

Iacopo Colonnelli is an Assistant Professor in the Department of Computer Science at the University of Turin, Italy. He is a member of the HiPEAC community, a member of the CINI HPC-KTT National Laboratory, serves on the Technical Committee of the Common Workflow Language (CWL), and is a founding coordinator of the CWL4HPC working group. He earned his Ph.D. in Modeling and Data Science with honors from the University of Turin, where his thesis on novel workflow models for heterogeneous distributed systems received the ITADATA 2023 Best PhD Thesis Award. He has co-authored over 40 peer-reviewed publications in national and international journals and conferences, and has contributed to more than 10 funded research projects. He is currently the local Principal Investigator for the DARE European project (total budget: €240M). He is also the designer and maintainer of the StreamFlow workflow manager (120+ citations on Google Scholar as of May 2026, recognized as an emerging technology by the EC Innovation Radar initiative) and has developed several other frameworks and libraries for workflow management and high-performance computing. His research interests include workflow modeling and management in heterogeneous distributed architectures, high-performance computing and I/O, distributed confidential computing, and large-scale data science.

Work experience

2023–present **Assistant Professor (RTD-A)**, *University of Turin*
Recipient of a 5-year contract as Assistant Professor at the Computer Science Department.

2022–2023 **Postdoctoral researcher**, *University of Turin*
Recipient of a research contract in the field of scientific workflow management systems, with a specific focus on distributed workflow modeling, in the context of the ACROSS European Project (G.A. 955648, total budget: €8M).

2021–2022 **Research engineer**, *University of Turin*
Recipient of a research scholarship in the field of scientific workflow management systems, with a specific focus on distributed workflow modeling.

2015–2018 **Technical team leader**, *Aktive Reply Srl*, Turin
Implementation of a Java-based learning management system on top of Liferay, a modular open source framework mainly based on the OSGi standard.

Honors and achievements

2026 **Global Young Scientists Summit**, *National University of Singapore*
Selected as one of two nominees from the University of Turin to participate in the 2026 Global Young Scientists Summit (GYSS) in Singapore — a prestigious international forum that brings together leading scientists and young researchers to discuss innovation, technology, and solutions to global challenges.

2023 **ITADATA 2023 Best PhD Thesis Award**, *CINI Data Science Lab*
The CINI Data Science Lab established a Distinguished Award for the Best PhD Thesis in all fields of Big Data and Data Science successfully defended at an Italian university in 2023. The theses have been evaluated by a committee of experts based on originality and potential impact on the Big Data and Data Science community. My PhD thesis, titled "Workflow models for heterogeneous distributed systems," has been selected as the best thesis among the 18 applications received.

2023 **EC Innovation Radar exploring technology award**, *EC Innovation Radar*
The Innovation Radar is an initiative of the European Commission to identify high-potential innovations and innovators in EU-funded research and innovation projects. In the context of the DeepHealth European project, the StreamFlow framework has been recognized as an "exploring technology."

2023 **Bando Premialità TO-PEOPLE 2023**, *University of Turin*
The University of Turin, in the context of the Three-Year Program 2021-2023 TO-PEOPLE, launched an initiative to recognize the skills and scientific productivity of young researchers at the early stage of their academic careers. I was among the winners of the 2023 edition.

2022 **FGCS Fall 2022 Editors' Choice**, *Future Generation Computer Systems editors*
Twice a year, the editorial board of the journal Future Generation Computer Systems selects a set of papers for the Editors' Choice. Selected papers have recently been published in the FGCS journal and are capturing the attention of readers. My article titled "Distributed workflows with Jupyter" has been included in the Fall 2022 Editors' Choice list.

2019 **HPC-Europa3 grant**, *HPC-Europa3 Scientific Users Selection Panel*
The HPC-Europa3 Transnational Access programme was a competitive initiative for European research scientists that provided access to HPC infrastructure across Europe. The program offered free access to HPC centers through a single application and an international two-stage peer-review process. The HPC-Europa3 Scientific Users Selection Panel evaluated applications based on scientific merit. I have been awarded an HPC-Europa3 grant for a 13-week visit to the Barcelona Supercomputing Center.

2019 **ACACES 2019 grant**, *High performance, edge and cloud computing (HiPEAC)*
I have been selected as a HiPEAC grant recipient for attending the 15th International Summer School on Advanced Computer Architecture and Compilation for High-performance Embedded Systems (ACACES 2019) in Fiuggi, Italy.

Education

2018–2022 **PhD in Modeling and Data Science**, *University of Turin*, Excellent with honors

2014–2017 **Master's Degree in Computer Engineering**, *Polytechnic University of Turin*, 110/110 cum laude

2011–2014 **Bachelor's Degree in Computer Engineering**, *Polytechnic University of Turin*, 110/110 cum laude

Research activity

2024–present **Founder and coordinator of the CWL4HPC working group**

The CWL4HPC working group aims to identify workflow patterns for modeling large-scale scientific applications and to implement the related Common Workflow Language (CWL) enhancement proposals. I am the leading proponent of the CWL Loop feature to model iterative workflow patterns, which has been accepted for inclusion in CWL v1.3.

2021–present **Member of the CINI HPC-KTT National Laboratory**

The CINI HPC Key Technologies and Tools (HPC-KTT) National Laboratory focuses on programming and execution models and tools, system software, high-performance software engineering, energy consumption reduction, and cloud computing. As a member of the Laboratory, I participated in four European projects funded by the EuroHPC Joint Undertaking: ACROSS (G.A. 955648, total budget: €8M), TEXTAROSSA (G.A. 956831, total budget: €6M), EUPEX (G.A. 101033975, total budget: €41M), and EUPILOT (G.A. 101034126, total budget: €30M).

2022–present **Member of the Common Workflow Language Technical Committee**

The Common Workflow Language (CWL) is an open standard for describing how to run command-line tools and connect them to create workflows. As a member of the technical committee, I am initially responsible for approving new versions of the CWL standards before the proposal goes to the CWL Leadership Team.

2022–present **Member of the Workflow Run RO-Crate working group**

Workflow Run RO-Crate is a working group for defining RO-Crate profiles for capturing the provenance of an execution of a computational workflow. Since 2022, I have been a member of the working group, and together we developed the first version of the Workflow Run RO-Crate (WRROC) profile collection, which led to a joint scientific publication on the PLOS One journal, titled "Recording provenance of workflow runs with RO-Crate."

2018–present **Member of the Alpha Parallel Computing Group**, *University of Turin*

The Parallel Computing research group (alpha) at the University of Torino, led by Prof. Marco Aldinucci, investigates and develops programming models, languages, and tools in parallel and distributed programming, high-performance computing, and federated learning. As a member of the group, I published 40+ research articles in peer-reviewed scientific journals and conferences and participated in 10+ national and international funded research projects. As an Assistant Professor, I am coordinating research on distributed workflows, distributed confidential computing, and low-level support for high-performance I/O.

2020 **Visiting Researcher at the Predictable Parallel Computing Group**, *Barcelona Supercomputing Center*

The Predictable Parallel Computing research group, led by Dr. Eduardo Quiñones Moreno, aims to converge the High-Performance Computing and critical real-time embedded systems domains. I participated in the group's activities under the HPC-Europa3 grant program for 13 weeks, developing methods and tools for high-performance data analysis within the DeepHealth European project (G.A. 825111, total budget: €14.8M).

Research projects

European projects

2025–present **Participant**, *IT4LIA*, Italy for Artificial Intelligence, EuroHPC JU, G.A. 101234224, total budget: €30M

2021–present **Participant**, *EUPILOT*, Pilot using independent, local and open technologies, EuroHPC JU, G.A. 101034126, total budget: €30M
Main achievement. We delivered the first publicly available RISC-V port of the PyTorch deep learning library.

2023–present **Task leader**, *SPACE CoE*, Centre of excellence on scalable parallel and distributed astrophysical codes for exascale, EuroHPC JU, G.A. 101093441, total budget: €8M
Main achievement. We developed the Cross-Facility Federated Learning (xFFL) methodology to train trillion-parameter models across multiple data centers.

2022–present **Participant**, *EUPEX*, European Pilot for Exascale, EuroHPC JU, G.A. 101033975, total budget: €41M
Main achievement. StreamFlow has been selected as one of the exploitable results of the project.

2024–present **Local principal investigator and task leader**, *DARE*, Digital autonomy for RISC-V in Europe, EuroHPC JU, G.A. 101202459, total budget: €240M

2021–2024 **Task leader**, *ACROSS*, HPC Big Data Artificial Intelligence cross-stack platform toward exascale, EuroHPC JU, G.A. 955648, total budget: €8M
Main achievement. StreamFlow plays a central role in the ACROSS software stack, allowing users to easily define application workflows and match them with the most suitable execution environment.

2021–2024 **Participant**, *TEXTAROSSA*, Towards extreme scale technologies and accelerators for EuroHPC HW/SW supercomputing applications for Exascale, EuroHPC JU, G.A. 956831, total budget: €6M

2019–2021 **Task leader**, *DeepHealth*, Deep-Learning and HPC to Boost Biomedical Applications for Health, Horizon 2020, G.A. 825111, total budget: €14.8M
Main achievement. StreamFlow has been recognized as an exploring technology by the EC Innovation Radar initiative.

National projects

2024–2025 **Participant**, *DEVOPS*, Prototype and integrate DevOps processes for ML model lifecycle (training, testing, deployment) into a multi cloud and multi-platform enterprise, ICSC Spoke 1 Innovation Grant program, total budget: €150k
In collaboration with Sogei SpA.

2023–2025 **Principal investigator**, *CPFWI*, Cross-Platform Full Waveform Inversion, ICSC Spoke 1 Innovation Grant program, total budget: €60k
In collaboration with ENI SpA.

2023–2025 **Work package and task leader**, *HaMMon*, Hazard mapping and vulnerability monitoring, ICSC Spoke 1 Innovation Grant program, total cost: €2M
In collaboration with Leithá Srl (Unipol Group).

2022–2025 **Participant**, *ICSC Spoke 1*, Future HPC and Big Data, Next Generation EU, total budget: €320M

Scientific conferences

Service as a chairperson

2026 **Posters and demo chair**, *Euro-Par 2026*, 32nd International European Conference on Parallel and Distributed Computing (Euro-Par 2026), Pisa, Italy

2025 **General chair**, *HiPES 2025*, 2nd Workshop about High-Performance e-Science, Dresden, Germany
Co-located with the 31st International European Conference on Parallel and Distributed Computing (Euro-Par 2025).

2025 **General and program chair**, *PDP 2025*, 33rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, Turin, Italy

2024 **General chair**, *WiDE 2024*, 2nd International Workshop on Workflows in Distributed Environments, Athens, Greece
Co-located with the 19th European Conference on Computer Systems (EuroSys 2024).

2024 **General chair**, *HiPES 2024*, 1st Workshop about High-Performance e-Science, Madrid, Spain
Co-located with the 30th International European Conference on Parallel and Distributed Computing (Euro-Par 2024).

2024 **General chair**, *CWLCon 2024*, 4th CWL Conference, Amsterdam, Netherlands

2023 **General chair**, *WiDE 2023*, 1st International Workshop on Workflows in Distributed Environments, Turin, Italy
Co-located with the 7th Annual Computers, Software, and Applications Conference (COMPSAC 2023).

2023 **Local chair**, *COMPSAC 2023*, 47th Annual Computers, Software, and Applications Conference, Turin, Italy

Committee memberships

2026 **Program committee**, *Middleware*, ACM/IFIP International Middleware Conference

2024–2025 **Program committee**, *WORKS*, Workshop on Workflows in Support of Large-Scale Science

2020, 2024–2025 **Program committee**, *HiPC*, IEEE International Conference on High Performance Computing, Data, and Analytics

2024–2025 **Program committee**, *ReWords*, Workshop on Reproducible Workflows, Data Management, and Security

2023–2025 **Program committee**, *SBAC-PAD*, International Symposium on Computer Architecture and High Performance Computing

2024–2025 **Program committee**, *SC*, International Conference for High Performance Computing, Networking, Storage, and Analysis

2024–2025 **Program committee**, *ICPP*, International Conference on Parallel Processing

2022–2024 **Program committee**, *HLPP*, International Symposium on High-level Parallel Programming and Applications

2024 **Program committee**, *IWAHPCE*, International Workshop on Arm-based HPC: Practice & Experience

2024 **Program committee**, *FRAME*, Workshop on Flexible Resource and Application Management on the Edge

2021–2023 **Program committee**, *Euro-Par*, International European Conference on Parallel and Distributed Computing

2022–2023 **Program committee**, *HPCMALL*, International Workshop on Malleability Techniques Applications in High-Performance Computing

2023 **Program committee**, *PDP*, Euromicro International Conference on Parallel, Distributed, and Network-Based Processing

2023 **Program committee**, *PAW-ATM*, Annual Parallel Applications Workshop, Alternatives To MPI+X

2019, 2022 **Program committee**, *IPTA*, International Conference on Image Processing Theory, Tools and Applications

Invited talks

2024 **Keynote speaker**, *WSCC 2024*, Scientific Workflows in the Continuum Era, Madrid, Spain

2024 **Invited speaker**, Cross-Facility Federated Learning - Part II, Helsinki, Finland

2024 **Invited speaker**, CWL in the HPC Ecosystem, CERN, Meyrin, Switzerland

2023 **Invited speaker**, *CWLCon 2023*, CWL for HPC: are we there yet?, EMBL, Heidelberg, Germany

2023 **Invited speaker**, Standardised Workflows at EBRAINS, Marseille, France

2023 **Invited speaker**, *OSA2Micro*, Workflows and the Common Workflow Language (CWL), Turin, Italy

2023 **Invited speaker**, *ITADATA 2023*, Workflow models for heterogeneous distributed systems, Naples, Italy

2022 **Invited speaker**, *2nd HealthyCloud Workshop*, Dossier: multi-tenant distributed Jupyter Notebooks, Virtual event

2022 **Invited speaker**, CINI HPC-KTT: HPC Key Technologies and Tools National Lab, Casalecchio di Reno, Italy

2021 **Invited speaker**, *CWLCon 2021*, The Universal Cloud-HPC Pipeline for the AI-Assisted Explainable Diagnosis of COVID-19 Pneumonia, Virtual event

Regular talks

- BookedSlurm: meeting user needs for advanced resource reservations in Slurm (eScience 2025) Chicago, IL, USA.
- Hybrid Workflows For Large-Scale Scientific Applications (EAGE 2022) Milan, Italy.
- 2023 CINI HPC-KTT National Assembly Pisa, Italy.
- Hybrid workflows for heterogeneous distributed computing (ITWSHPC 2022) Torino, Italy.
- JupyterFlow: Jupyter Notebooks su larga scala Rome, Italy.
- J on The Beach 2022 Malaga, Spain.
- Workshop GARR 2019 Rome, Italy.
- A Systematic Mapping Study of Italian Research on Workflows (WORKS 2023) Denver, CO, USA.
- Cross-Facility Federated Learning (EuroHPC 2023) Bruxelles, Belgium.

- 2025 ICSC Spoke 1 Event Milan, Italy.
- 2023 Workflows Community BoF: Modern Workflows for Continuum and Cross-Facility Computing Denver, CO, USA.
- Towards a European AI Platform (EuroHPC 2024) Amsterdam, Netherlands.
- CWL Working Groups (CWLCon 2024) Amsterdam, Netherlands.
- Deep Learning at Scale (PDP 2019) Pavia, Italy.

Open source software

2024–present **SWIRL**, *Designer and maintainer*

SWIRL is an intermediate representation language for distributed and decentralized scientific workflows. The `swirlc` compiler can be used to translate a traditional workflow instance (written in Galaxy, CWL, or Pegasus) into a low-level, fully distributed executable targeting edge and fog computing environments.

2020–present **CAPIO**, *Developer and maintainer*

CAPIO is a middleware that injects I/O streaming capabilities into file-based workflows, improving the computation-I/O overlap without requiring changes to the application code. Until now, it has served five funded research projects and is part of several scientific collaborations with academic and industrial partners (INAF, Eni SpA, Morfo Srl, INGV).

2021–present **Jupyter Workflow**, *Designer and maintainer*

Jupyter Workflow is a JupyterLab extension designed to support distributed literate workflows and to execute them on hybrid cloud/HPC infrastructures. It is now part of the DARE Project's software stack. Its associated publication in the Future Generation Computer Systems journal, titled "Distributed workflows with Jupyter," has been included in the "FGCS Fall 2022 Editors' Choice" list.

2019–present **StreamFlow**, *Designer and maintainer*

StreamFlow is a container-native Workflow Management System based on the Common Workflow Language standard that supports hybrid workflows and their execution across cloud/HPC infrastructures. Until now, it has served eight National and International research projects and is part of several scientific collaborations with academic and industrial partners (INAF, INGV, EBRAINS, ASTRON, IBM, Sogei SpA, Eni SpA). It was recognized as an "emerging technology" by the EC Innovation Radar initiative in 2023 and selected as one of the three exploitable results of the EUPEX project. Its associated publication in the IEEE Transactions on Emerging Topics in Computing journal, titled "StreamFlow: cross-breeding cloud with HPC," counts 120+ citations on Google Scholar.

Editorial activity

2025–present **Subject area editor**, *Parallel Computing*, Elsevier

2024–2025 **Guest editor**, *Future Generation Computer Systems*, Elsevier

Special Issue on High-Performance eScience key technologies, tools, and applications

Review activity

Research projects

2025 **SNSF Ambizione**

Reviewer of a project proposal under the "2025 SNSF Ambizione" funding program of the Swiss National Science Foundation (SNSF).

2024 **EuroHPC Joint Undertaking**
Reviewer of the "An Optimization and Co-design Framework for Sparse Computation" (SparCity) European project (G.A. 956213, total budget: €2.6M).

2023 **2023 CINECA ISCRA program**
Reviewer of two project proposals under the "Italian SuperComputing Resource Allocation" (ISCRA) initiative of the CINECA consortium.

2022 **2022 IFAB Call for Projects**
Reviewer of three project proposals under the "2022 IFAB Call for Projects" funding program of the International Foundation Big Data and Artificial Intelligence for Human Development (IFAB).

Scientific journals

- BMC Bioinformatics, Springer, ISSN 1471-2105.
- Future Generation Computer Systems, Elsevier, ISSN 0167-739X.
- F1000Research, ISSN 2046-1402.
- IEEE Transactions on Network and Service Management, IEEE, ISSN 1932-4537.
- International Journal of Parallel Programming, Springer, ISSN 1573-7640.
- International Journal of Intelligent Systems, Hindawi, ISSN 1098-111X.
- Journal of Parallel and Distributed Computing, Elsevier, ISSN 1096-0848.
- Nature Scientific Reports, Springer, ISSN 2045-2322.
- Parallel Computing, Elsevier, ISSN 0167-8191.
- SoftwareX, Elsevier, ISSN 2352-7110.
- Software and Systems Modeling, Springer, ISSN 1619-1374.
- Scientific Programming, Hindawi, ISSN 1875-919X.
- The Journal of Supercomputing, Springer, ISSN 1573-0484.

Teaching activity

Regular courses

2023/2024–2025/2026 **Lecturer**, *Computer Science*, University of Turin, Drug Science and Technology Dept., Undergraduate course
16h of lectures, 16h of exercises.

2023/2024 **Lecturer**, *Methods and tools for HPC and convergent computing*, University of Turin, Computer Science Dept., PhD course
Module "Cloud computing," 4h of lectures, 4h of hands-on tutorial.

Professional Master's programs

2025–2026 **Invited lecturer**, *Master in High-Performance Computing*, SISSA, Trieste, Italy
Module "Hybrid workflows for large-scale AI," 10h of lectures, 10h of hands-on tutorial.

2023/2024 **Invited lecturer**, *Medicina di Precisione tra cultura e società: Genomica, Bioinformatica, AI e Management*, Università Campus Bio-Medico di Roma, Rome, Italy
Module "High-performance computing," 8h of lectures, 8h of hands-on tutorial.
Supervision of a student's final project titled "SNV Analysis Workflow in CWL, with hybrid deployment on HPC and local using StreamFlow."

2024 **Invited lecturer**, *CUAP Data Steward for Open Science*, University of Turin, Italy
Module "Software environments for research: reproducibility, virtual environments, workflows," 8h of lectures.

2022/2023 **Invited lecturer**, *Cloud Computing per Intelligenza Artificiale e Internet of Things*, University of Turin, Italy
Module "DevOps & Automation," 8 hours lecturing, 2 hours of hands-on tutorial.
Module "Containerization platforms & Virtualization," 18h of lectures, 18h of hands-on tutorial.

PhD summer/winter schools

2025 **Invited lecturer**, *Advanced Topics in Machine Learning: Machine Learning Operations MLOps*, Technical University of Denmark, Copenhagen
Module "Federated Learning: a gentle introduction to federated optimization," 1.5h of lectures, 1.5h of hands-on tutorial.

2024 **Invited lecturer**, *Advanced School on Applied Machine Learning*, ICTP, Trieste, Italy
Module "Dynamic hybrid workflows for deep learning on high-performance computing infrastructures," 1.5h of lectures.

2023 **Teaching assistant**, *National Summer School on High Performance Computing*, University of Pavia, Italy
Module "Workflows for future high-performance-computing," 1.5h of lectures, 1.5h of hands-on tutorial.

2022 **Teaching assistant**, *DeepHealth Winter School*, University of Turin, Italy
Module "High-performance computing," 1.5h of hands-on tutorial.

Mentorships

2023–present **Co-supervisor**, *PhD student*, End-to-end confidential workflows (working title), University of Turin

2023–present **Co-supervisor**, *PhD student*, Large-scale workflow orchestration (working title), University of Turin

2025/2026 **Co-supervisor**, *Bachelor's thesis*, Cloud-Oriented Workflow Scalability, University of Turin

2024/2025 **Supervisor**, *Bachelor's thesis*, Transparent user-space networking, University of Turin

2024/2025 **Co-supervisor**, *Master's thesis*, Distributed workflows in Rust, University of Turin

2023/2024 **Co-supervisor**, *Bachelor's thesis*, Human–machine interface for visualization, debugging, and editing of CWL workflows, University of Turin

2023/2024 **Supervisor**, *Professional Master's final project*, SNV Analysis Workflow in CWL, with hybrid deployment on HPC and local using StreamFlow, Università Campus Bio-Medico di Roma

2023/2024 **Supervisor**, *Bachelor's thesis*, Syscall Intercept: RISC-V porting, University of Turin

2023/2024 **Supervisor**, *Bachelor's thesis*, Human–machine interface for visualization, debugging, and editing of HPC workflows, University of Turin

2022/2023 **Co-supervisor**, *Master's thesis*, Trusted computation for the execution of distributed workflows, University of Turin

2022/2023 **Co-supervisor**, *Master's thesis*, Porting the variant calling pipeline for NGS data in cloud-HPC environment, University of Turin