

CIDAI Centre of Innovation for Data tech and Artificial Intelligence

Personalized Evacuation Routing System: A Smart Solution for Crisis Management

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Generalitat de Catalunya



CIDAI-POC-2025-05
Some rights reserved. This work is available under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0).

Under the terms of this licence, you may copy, redistribute and adapt the work for non-commercial purposes, provided the work is appropriately cited, as indicated below.
In any use of this work, there should be no suggestion that CIDAI endorses any specific organization, products or services. The use of the CIDAI logo is not permitted.
If you adapt the work, then you must license your work under the same or equivalent Creative Commons licence.

If you create a translation of this work, you should add the following disclaimer along with the suggested citation: “This translation was not created by the Centre of Innovation for Data tech and Artificial Intelligence (CIDAI). CIDAI is not responsible for the content or accuracy of this translation. The original English edition shall be the binding and authentic edition”.
Any mediation relating to disputes arising under the licence shall be conducted in accordance with the mediation rules of the World Intellectual Property Organization.

Suggested citation. CIDAI-POC-2025-05 // Personalized Evacuation Routing System: A Smart Solution for Crisis Management by Eduardo Quiñones, Xavier Palomo and Janine Gehrig – Barcelona Supercomputing Center, CIDAI, 2025. Licence: CC BY-NC-SA 4.0



I.	ABSTRACT	4
II.	DESCRIPTION OF THE PROBLEM SOLVED	7
III.	IMPLEMENTATION OF THE SOLUTION	9
IV.	IMPACT OF THE SOLUTION	13
V.	ACKNOWLEDGMENTS	14
VI.	CITED WORKS	15



I. Abstract

A. Summary

With more than 180 zB of data worldwide, cities have more opportunities than ever for pursuing data-driven policies that can help their citizens face pressing urban challenges. Technologies like AI, edge, cloud, HPC and more can use these extreme amounts of data for the benefit of the economy and society. However, there are currently barriers to accessing and using this technology.

To bridge the gap between data availability and knowledge for making informed decisions, BSC is leading the EXTRACT project, which is developing a distributed data-mining software platform for extreme data across the continuum. This platform seamlessly integrates different cutting-edge computing environments (edge, cloud, and HPC) to ensure that decision-makers are able to access the right knowledge and compute resources at the right time.

Through its Personalized Evacuation Routing System (PER) use case, implemented in Venice, Italy, the EXTRACT project is showing the viability of the project technology for extracting valuable and trustworthy knowledge from the huge and varied volumes of data that exist and convert it into actionable knowledge (instructions to a smartphone) that can optimize decision-making in times of crisis.

The project targets the entire lifecycle of extreme data mining workflows. In the case of PER, this means unifying data from European services like Copernicus and Galileo, 5G localization signals, and IoT sensors embedded in smart city infrastructures to help provide support for crisis management. The EXTRACT project harnesses the power of data and tests its technology to guide citizens in an urban

environment through a safe route in real time. This use case demonstrates the value of EXTRACT technologies for addressing challenges with high geographical and social complexity.

EXTRACT technology was successfully tested in the Venice PER use case and a Minimum Viable Product (MVP) was presented in September 2024 is currently being further refined to include greater functionalities. EXTRACT technology is also being used in the case of urban mobility and pollution in the city of Barcelona through the ASCENDER project.

B. Objectives

To address key challenges posed by extreme data and to check its usefulness for key activities like urban crisis management, the EXTRACT project is:

- Developing novel and more holistic approaches that combine cutting edge HPC, edge and cloud computing technologies into a unified solution that is able to securely manage and analyse the complete lifecycle of extreme data, effectively addressing its specific characteristics.
- Creating new development and execution platforms to easily describe data mining workflows deployed across these different computing environments.
- Testing the usefulness of this technology in two use cases with different extreme data requirements: the Personalized Evacuation Routing (PER) system urban crisis management use case, discussed in this document, and the Transient Astrophysics with a Square Kilometre Array Pathfinder (TASKA) astronomy use case.

C. Challenge that is being addressed

To take advantage of the opportunities offered by extreme amounts of data for decision-making, a combination of cutting-edge technologies is needed to handle the complete data lifecycle and value chain. Huge volumes of data need to be collected across distributed and heterogeneous sources. This data needs to be mined for meaningful, accurate, reliable and useful knowledge, and applications and end users must be able to consume the data in a secure and trustworthy manner.

Current solutions start to fail when coping with extreme data (large volume, complexity, variety) because programming practices and computing technologies are only optimized to cope with some data characteristics, rather than addressing them in a holistic manner. For example, while *High-Performance Computing (HPC) technologies* can effectively address volume and speed characteristics and can enable the training and execution of complex AI, they are not suitable for real-time operations, geographically dispersed data sources, extreme large-stored volumes, and they are not energy-efficient. *Edge computing* moves computation closer to data sources where data is captured, reducing data transmission latencies while minimizing data privacy and security risks, reducing energy consumption. However, it lacks the computing power needed to execute complex AI methods addressing complexity and incompleteness, store large volumes and provide high-assurance security and privacy primitives. *Cloud computing* provides highly scalable storage systems however, it suffers from communication latencies and resource contention.

In the case of urban areas that may face crisis scenarios like earthquakes, fires, floods, acts of terrorism and more, emergency services need to quickly and confidently use data to make decisions without thinking about the technology and processes behind it. Disasters trigger sudden and large data spikes that can be considered extreme in terms of volume, heterogeneity and their dynamics. Processing and analysing this data in real time is crucial for implementing evacuation strategies.

As decision-makers become increasingly aware of the importance of preparedness, advanced technology is proving to be a critical tool for turning vast amounts of urban data into actionable insights. This is particularly true for cities like Venice, whose geographical complexity presents specific logistical challenges. The city of Venice is composed of 118 small islands that are separated by canals and linked by over 400 bridges. Moreover, the historical area of Venice is usually overcrowded with unprepared tourists walking in an unknown city that has 95% of its territory at flooding risk. The intricate network of narrow lanes and canals in Venice makes traditional evacuation strategies inadequate. Protecting both residents and the millions of tourists who visit Venice yearly is a real challenge. The PER use case provides a perfect testing ground for evaluating the usefulness of EXTRACT technology

II. Description of the problem solved

The personalized evacuation routing in the EXTRACT use case represents a major leap forward in urban emergency management. Instead of issuing generic, citywide directives, the system produces evacuation instructions adapted to each individual's real-time context. This strategy emphasizes preparedness, flexibility, and adaptability—core principles that reinforce the community's overall resilience. By tailoring routes and instructions, it is possible to adapt to changing disaster conditions in real-time to ensure that response efforts evolve seamlessly with the crisis and feed into a cycle of continuous improvement.

The EXTRACT solution plays a key role here by collecting, analysing, and fully utilizing the rich streams of "extreme" data available to the city of Venice. Thanks to a wide range of sources—spanning urban

Figure 1: Infographic representation of the PER use case in Venice.





infrastructure data, real-time behavioural inputs, edge IoT sensors, 5G networks, simulated disaster scenarios, as well as positioning and mapping from Galileo and Copernicus—city administrators gain an avenue for designing data-intensive crisis management solutions. By effectively processing and merging these disparate data sets, we not only reduce response times but also develop far more refined, context-aware evacuation protocols.

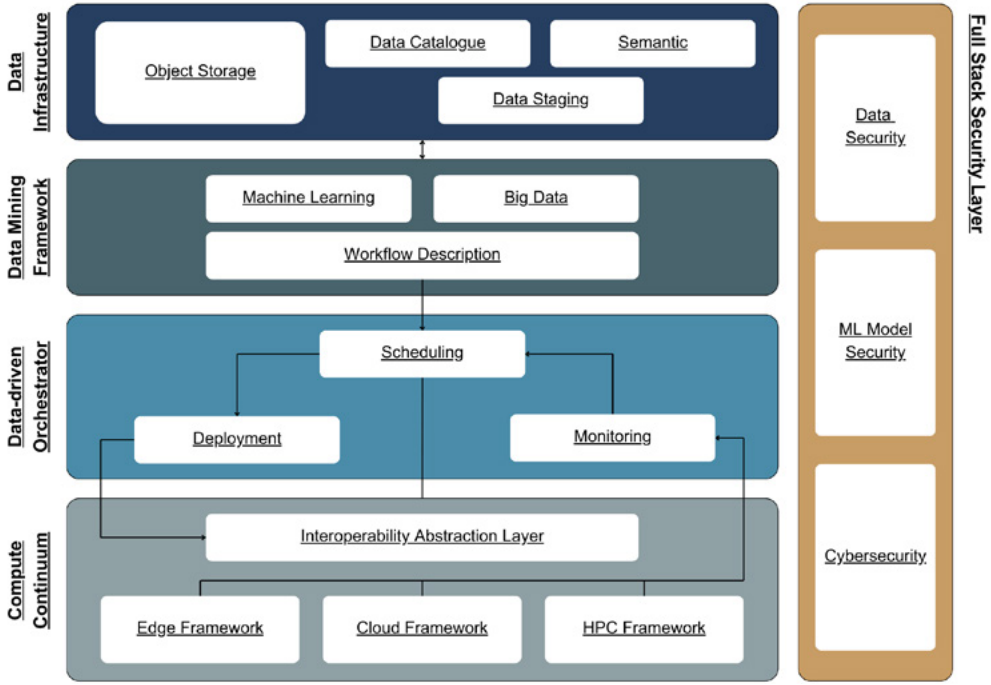
Concretely, in the PER use case, we harness these data flows to guide citizens through safe evacuation routes in near real time. This approach exemplifies how EXTRACT technology addresses the complicated interplay of geographical and social elements found in a city like Venice. The platform enables the rapid assimilation, processing, and semantic annotation of huge amounts of highly varied data—an essential ingredient in agile and efficient decision-making. By delivering personalized, immediate guidance to each citizen on where to take shelter, the PER system diminishes vulnerability and bolsters the community's sense of security. It also supports first responders, helping them coordinate evacuations in a way that avoids congesting certain routes, overcrowding select shelters, or sowing confusion with inconsistent messages. Ultimately, PER aspires to be a transformative resource in urban crisis response, combining leading-edge technologies and robust data pipelines to create a singular system that thrives on the volume and quality of incoming signals and models.

III. Implementation of the solution

A. Architecture

The EXTRACT platform provides a cohesive software stack that connects edge nodes, cloud resources, and HPC systems, creating what is effectively a single, unified environment.

Figure 2: EXTRACT software stack



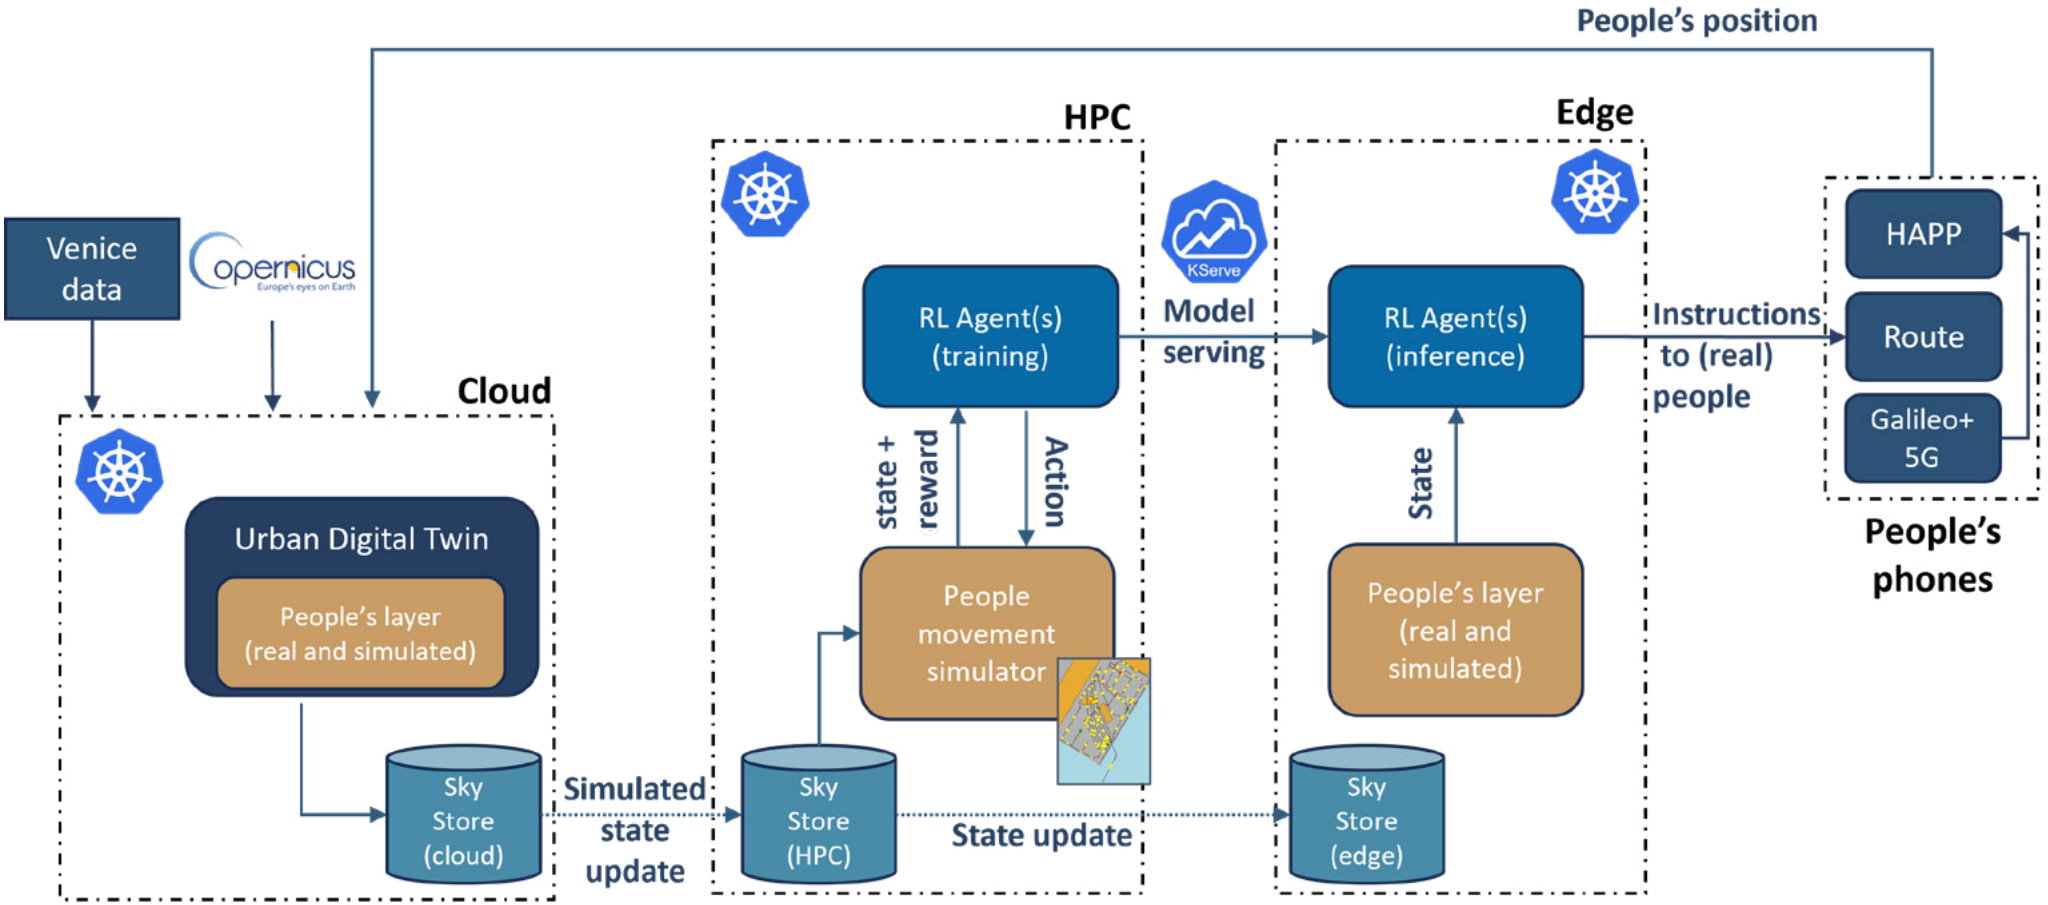
In the context of this project, an application typically starts by gathering data—whether from sensor logs at the edge, massive data collections in the cloud, or simulation outputs in HPC clusters. That data flows into the Data Infrastructure layer, which organizes and annotates the information, ensuring that every resource (from streaming sensor updates to HPC archives) can be accessed in a consistent, transparent manner.

Once the data is staged and ready, the Data Mining Framework takes over, enabling developers or experts to define analytical workflows. Tools like COMPSs or Lithops help break these workflows into discrete tasks that can be deployed to the most suitable compute node—edge, cloud, or HPC—depending on needs like performance or real-time constraints. At the same time, the Data-Driven Orchestrator acts as the operational brain of EXTRACT, continuously monitoring resource load, network latencies, or cost parameters, and allocating or redeploying tasks on the fly. Meanwhile, HPC clusters, cloud backends, and edge nodes are grouped under the Compute Continuum layer, letting an application seamlessly tap into whichever environment is optimal. A robust Full Stack Security layer spans all these components, safeguarding data transfers, model training, and overall system continuity against security threats.

B. Implementation

For the PER scenario in Venice, the EXTRACT platform converges a series of powerful components to produce a robust, real-time guidance system that addresses the city's complex topography and fluctuating crowd dynamics.

Figure 3: PER infrastructure diagram.



At the core is the Urban Digital Twin (UDT), which merges historical architectural data, high-resolution maps, and up-to-date human activity. This digital "mirror" of Venice not only offers a snapshot of current conditions (such as open streets, blocked canals, or crowd movements) but also supports predictive modeling. In practice, it compiles input from IoT sensors, 5G signals, and geolocation data, enabling the city's administrators and first responders to simulate potential emergencies on a virtual yet highly realistic canvas.

Building on the UDT, the Multi-Agent Reinforcement Learning (MARL) system evaluates and continuously refines potential evacuation pathways. Each "agent" in the system is programmed to move within the digital twin as a stand-in for real citizens, adjusting its route based on factors like local congestion or newly emerging hazards (e.g., a flooded district or a structural collapse). By running repeated training cycles, the MARL approach yields intelligent evacuation strategies capable of adapting to new configurations or critical events in real time, thus significantly reducing bottlenecks and confusion when an actual emergency arises.

To thoroughly test these strategies, the project relies on a dedicated Simulator capable of running massive "what-if" scenarios. This environment ingests data flows from the UDT and MARL, reconstructing how individuals would move and respond to a crisis. For instance, it might simulate thousands of tourists stuck near certain bridges or local residents hurrying to designated safe zones. Each test reveals areas for improvement, shaping more refined, safer evacuation plans over repeated cycles.

Underpinning it all is a robust pipeline for collecting, managing, and preprocessing data, this is the platform architecture that connects all the required technologies and provides an interoperability layer. Inputs from the Smart Control Room (SCR), HPC nodes, or cloud-based systems feed into this framework, ensuring timely and uniform access to essential statistics. For the PER use case, this includes near-real-time crowd density readings, updated hazard zones, or even external data from local networks. By providing a single, coherent view of the city's operational data, the framework ensures the entire pipeline—from MARL computations to HPC-based simulations—runs smoothly with minimal overhead.

Although the ultimate goal is a fully-fledged system that seamlessly delivers personalized escape routes to end-user devices, the project team first delivered an MVP focused on these backend functionalities. Released in September 2024, this MVP demonstrated crucial features:

- Data Handling: The ability to retrieve, merge, and process large, disparate data sets in near real time.
- Simulations and Route Optimization: Early integration of the UDT, MARL, and the simulation environment to produce basic, albeit partial, evacuation route solutions. For the MVP, we did not have stringent real time constraints, and only a small part of the city of Venice was considered in the map.

However, the MVP is not an endpoint. It merely confirmed that the EXTRACT platform can cope with "extreme data" and produce feasible escape routing strategies; many user-facing components—such as the Evacuation Mobile App (EMA) and real-time push noti-



fications—remain under development. Ongoing phases now focus on bridging the gap between these backend proofs-of-concept and a comprehensive, citizen-ready solution that not only offers robust data-driven guidance but does so through accessible, responsive interfaces.

Ultimately, the PER scenario exemplifies how a well-structured HPC–cloud–edge data pipeline, reinforced by advanced modeling and simulation, can tackle inherently chaotic urban emergencies. By refining each layer of the EXTRACT platform, the project paves the way toward a thoroughly integrated system where real-time crises prompt an equally dynamic, well-coordinated response, thereby enhancing both citizen safety and municipal resilience.

More details on the MVP can be found in EXTRACT deliverable D1.2 First release of the EXTRACT Use Cases

IV. Impact of the solution

The PER acts as a testbed for validating and demonstrating the platform's capabilities, but the EXTRACT software platform is broadly applicable, as shown by the two different use cases. Any scenario requiring distributed data processing, analytics pipelines, or heterogeneous resource orchestration can benefit from EXTRACT's cohesive approach. By merging comprehensive data management, a robust data mining framework, sophisticated orchestration, an interoperable compute continuum, and strong security, the platform stands ready to address a wide variety of domains far beyond its initial industrial pilots.

The EXTRACT project integrates advanced technologies that unlock new possibilities across a wide range of applications. It leverages tools such as COMPSS and Lithops to enable applications to harness parallelism and seamlessly operate across heterogeneous systems composed of different resources, such as edge, cloud or HPC, while maintaining the illusion of a unified, local system that executes locally for the user.

Beyond computing efficiency, EXTRACT technology also transforms access to an interaction with data. Whether data is stored nearby or distributed across multiple locations, the system abstracts this complexity, allowing users to execute highly demanding tasks without being constrained by data distribution challenges. This is crucial for applications that require real-time processing under stringent deadlines, such as in the case of the PER, where vast amounts of data need to be ingested and analysed quickly.

The true potential of EXTRACT technology lies in its ability to derive valuable and trustworthy insights from the huge volume of heterogeneous and dispersed raw data. By efficiently capturing, processing and integrating scattered raw data, it empowers data-driven decision-making across various domains, such as societal, scientific, environmental, economic, ones, thus paving the way for innovative solutions with far-reaching impacts.





V. Acknowledgments

The EXTRACT Project has received funding from the European Union's Horizon Europe programme under grant agreement number 101093110.

Special thanks to LOGOS-RI, Mathema SRL and the Municipality of the City of Venice.

Índex

VI. Cited works

Arjona, A., et al. Dataplug: Unlocking Extreme Data Analytics with On-the-fly Dynamic Partitioning of Unstructured Data. Zenodo, 2024, doi:10.5281/zenodo.14161823.

Arjona, A., et al. Exploring Secure and Efficient Temporary Data Sharing Between Co-located Kubernetes Containers. Zenodo, 2025, doi:10.5281/zenodo.15097005.

Arjona, A., et al. Scaling a Variant Calling Genomics Pipeline with FaaS. Association for Computing Machinery, 2023, New York, NY, USA, 59–64. doi: 10.1145/3631295.3631403.

Barcelona-Pons, D., et al. Glider: Serverless Ephemeral Stateful Near-Data Computation. Association for Computing Machinery, 2023, New York, NY, USA, 247–260. doi: 10.1145/3590140.3629119.

Bellini, A., et al. Multi-party Computation for Privacy and Security in Machine Learning: a Practical Review. 2023 IEEE International Conference on Cyber Security and Resilience (CSR). 2023, Venice, Italy, doi:10.1109/CSR57506.2023.1022482.

Cañete-Garrucho, A., The EXTRACT Project: Abstracting a highly distributed and heterogeneous infrastructure. Oral presentation, HiPEAC workshop, Advancing the compute continuum: intelligent solutions for extreme data and 5G-enabled ecosystems, 21 January 2025, Barcelona, Spain.

Eizaguirre, G. T.y M. Sánchez-Artigas. A Seer Knows Best: Auto-Tuned Object Storage Shuffling for Serverless Analytics. Elsevier, 2024, doi:10.1016/j.jpdc.2023.104763.

Eizaguirre, G. T.y M. Sánchez-Artigas. Is Performance of Object Storage Predictable for Serverless I/O Workloads? A Comparative Study. IEEE Computer Society, 2023, doi:10.1109/ICNP59255.2023.10355617.

Eizaguirre, G. T., et al. Serverful Functions: Leveraging Servers in Complex Serverless Workflows (Industry Track). Association for Computing Machinery, 2024, doi:10.1145/3700824.3701095.

EXTRACT. D1.1 Use-case Requirements. 1.0, Zenodo, June 30th, 2023, doi:10.5281/zenodo.15126401.

EXTRACT. D1.2 First Release of the EXTRACT Use-cases. Zenodo, February 29th, 2024, doi:10.5281/zenodo.15132002.

Gimeno Sarroca, P. MLLess: Achieving Cost Efficiency in Serverless Machine Learning Training. Journal of Parallel and Distributed Computing, vol 183, doi:10.1016/j.jpdc.2023.104764.

CIDAI Centre of Innovation for Data tech and Artificial Intelligence

C/Bilbao, 72 Edif. A. 08005 - Barcelona
Tel. +34 93 7419 100
info@cidai-catalonia-ai.eu
www.cidai.eu

Socis

Generalitat de Catalunya
Ajuntament de Barcelona
Barcelona Activa
BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación
CVC Computer Vision Center
eurecat Centre Tecnològic de Catalunya
i2cat
HUAWEI
Microsoft
NTT DATA
SAP
sdg group
UPC Universitat Politècnica de Catalunya BarcelonaTech Intelligent Data Science and Artificial Intelligence Research Center
IDEAI

Membres

3cat
ATM Àrea de Barcelona Autoritat del Transport Metropolità
CETAQUA WATER TECHNOLOGY CENTRE
AMB
Grup Caixa Enginyers Banca Cooperativa
CRUILLA