



Consiglio Nazionale
delle Ricerche



Istituto di Scienza e Tecnologie
dell'Informazione "A. Faedo"

ISTI News

The biannual magazine of the Institute of Information Science and Technologies "A. Faedo"

ISSUE 16

December 2024

Cover Story
**Generative AI For
Images and Virtual
Worlds**

EDITORIAL

EC enlightens a new digital
ecosystem for museums and
Cultural Heritage professionals

ISSN 2724-6566 DOI: 10.32079/ISTINews © Institute of Information Science and Technologies "Alessandro Faedo"

Contents

Editorial

EC Provides a New Digital Ecosystem for Museums and Cultural Heritage Professionals 3

Cover Story

Generative AI for Images and Virtual Worlds. 5

New Projects

MMinE-SwEEPER 9

MUSIT 10

Selected Papers

Indoor Localization Algorithms Based on Angle of Arrival with a Benchmark Comparison. 11

What are Data Spaces? Systematic Survey and Future Outlook 11

Formal Methods in Industry. 12

The Role of Formal Methods in Computer Science Education 13

Advancing Orchestration Synthesis for Contract Automata . 14

RETORCH*: A Cost and Resource Aware Model for E2E Testing in the Cloud. 15

A Risk Assessment Framework based on Fuzzy Logic for Automotive Systems 15

AI ‘News’ Content Farms are Easy to Make and Hard to Detect: a Case Study in Italian 16

Using Large Language Models to Create Narrative Events. 17

A Semantic Knowledge Graph of European Mountain Value Chains 17

A Conversational Agent for Creating Automations Exploiting Large Language Models. 18

Extracting Mediterranean Hidden Fishing Hotspots Through Big Data Mining 19

Climate Change Effects on Animal Presence in the Massaciuccoli Lake Basin. 20

Fast, Interpretable, and Deterministic Time Series Classification with a Bag-of-Receptive-Fields. 21

Describing Group Evolution in Temporal Data Using Multi-Faceted Events 22

Dimension Importance Estimation for Dense Information Retrieval. 23

Efficient Inverted Indexes for Approximate Retrieval over

Learned Sparse Representations 23

Adaptive Machine Learning Approach for Importance Harnessing Topological Machine Learning in Raman Spectroscopy: Perspectives for Alzheimer’s Disease Detection via Cerebrospinal Fluid Analysis 24

Environmental Study and Stress-Related Biomarkers Modifications in a Crew During Analog Astronaut Mission EMMPOL 6 24

ANN Uncertainty Estimates in Assessing Fatty Liver Content from Ultrasound Data 25

Alignable Lamella Gridshells. 26

Capacitive Touch Sensing on General 3D Surfaces. 27

Seismic Response and Ambient Vibrations of a Medieval Tower in the Mugello Area (Italy) 28

Software&Tools

A Configurable Software Model of a Self-Adaptive Robotic System. 29

Awards and Achievements

Best Paper Runner-up Award at SIGIR 2024 30

Best Paper Honorable Mention at ACM SIGGRAPH 2024 . 30

Best Talk Award at CCS2024 31

Best Paper Award at ASONAM 2024. 31

Cor Baayen Young Researcher Award 2024 32

Best Paper Award: Individual Alternative Routing based on Road Popularity 33

“Leonardo Lesmo” 2024 Award 33

Best Poster Award, PhD Day 2024 of DII 34

Grants for Young Mobility 35

Miscellaneous

Indoor Localization Technologies for Cultural Heritage: the MVL app 36

AI/ML Assisted 5G-NTN Integration for Optimized Network Slicing over Multi-Technology Architecture 36

Marco Malvaldi Joins ISTI to Unlock the Potentials of LLMs for Writing. 37

SoBigData Academy 38

Ph.D. Dissertations 40

Welcome aboard. 42

Conferences - Co-organized by ISTI 43

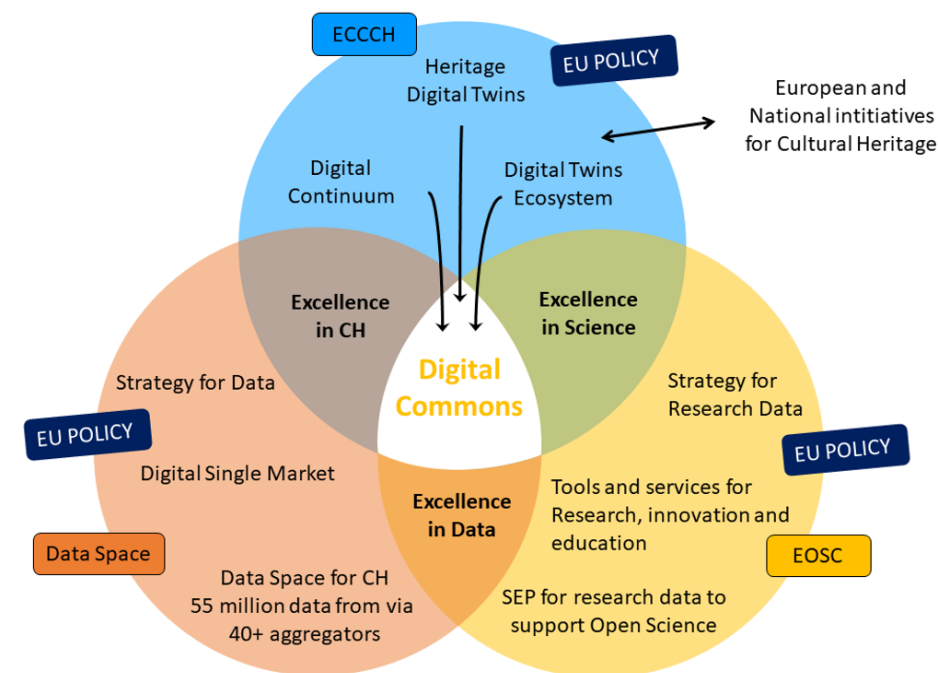
EC Provides a New Digital Ecosystem for Museums and Cultural Heritage Professionals

What will museums look like in the near future? How will Cultural Heritage (CH) professionals collaborate efficiently and ensure that the results of their work are properly maintained and accessed?

Digital technologies will increasingly play a role, following a trend that several leading institutions have already pioneered. The old 19th-century model, with numerous objects displayed with only tiny labels and a few written panels, will be transformed with more specific contexts. Collections will no longer be just for expert consumption, but will be informative and participatory. Fun environments will convey the stories of both the works and our past, aided by visual digital technologies and relying on mobile digital tools (increasingly using our phones/tablets, equipped with specific apps, or VR/XR headsets).

We need instruments that are easy to use which tell the story of the artworks and connect them to the culture of a community. This evolution could be monumental and unstoppable (just look at how storytelling has already changed in documentaries or in historical or fantasy films). However all this is currently hindered by the very small size of the vast majority of cultural institutions, the scarcity of financial and human resources, and the fact that these institutions often work in isolation, thus deploying insufficient financial resources to attract private enterprises.

These shortcomings have been recognized by the European Commission, which aims to give a strong boost to this sector by allocating approximately 100 million euros to finance research for building a specific digital and cooperative platform. In contrast with the more common autarchic and isolated mode, cooperation between institutions is the key to boosting the sector. The aim was to provide a common context where both the community of European museums



and the world of academia and conservation would exchange data, experiences and cooperate to set up a new era of museum communication. This is of course with the understanding that the museums of the future are not only about exhibitions and communication with the public, but also about research, conservation, and restoration.

From 2021 to 2022 the EC commissioned a select group of eight experts, to design a plan to achieve these goals. The director of ISTI had the honor of being part of this committee, which included specialists in museum management, communication, and digital technologies. This preparatory work resulted in a comprehensive report¹, which

1 P. Brunet, L. De Luca, E. Hyvönen, A. Joffres, P. Plassmeyer, M. Pronk, R. Scopigno, G. Sonkoly, "Report on a European collaborative cloud for cultural heritage: ex - ante impact assessment", prepared for European Commission Directorate-General for Research and Innovation, 2022. <https://op.europa.eu/en/publication-detail/-/publication/90f1ee85-ca88-11ec-b6f4-01aa75ed71a1>

served as the foundation for defining the vision, structure, and organization of subsequent EC calls.

The European vision was therefore not only to build a common database, which is a purely data oriented approach which has generally had little success despite being attempted many times, including in Italy. Instead, the aim is to launch a large integrated program to build a digital ecosystem composed of operational tools (practical digital applications) working on a common knowledge system. The overall aim is to support the collaborative activities of museums and professionals. These applications are intended to cover the entire range from the digitization of artworks to virtual or assisted museum navigation applications, tools for creating interactive installations, instruments facilitating the assisted design of new museum spaces, monitoring and controlling visitor flows, enabling the easy production of storytelling-based contents, applications to support the study of works (digital twins) including conservation and restoration issues, the introduction of supportive AI technologies, and the management of intellectual property rights (IPR).

This vision has been implemented with the European Collaborative Cloud for Cultural Heritage (ECCCH) initiative, which goes beyond the old European financing approach of multiple independent research projects, whose results have often been lost at the end of the average three-year project lifespan. In this new model, the project consortia selected after each call are asked to work on a common platform, which has to provide a set of shared tools and components to ensure: the use and enrichment of a common repository (storing data and metadata), a good degree of uniformity in the applications created (common interface design policies and common visualization and data access tools), a single common front-end to provide access to applications and results, and tools for evaluating the quality and effectiveness of those applications.

The idea is to establish a virtuous cycle of maintenance and future evolution of the most widely used results, following an impact evaluation provided by assessment data produced by the community of users. Within the framework of EC initiatives for the Cultural Heritage, this represents the first attempt to lay the foundations for an effective assessment and preservation policy, paving the way for the selective maintenance of the tools resulting from the research.

This European vision has taken shape through a series of ECCCH project calls (2023-2028), in which the first call (with a budget of €25 million) was awarded to a consortium coordinated by the French CNRS, with a considerable contribution from the Italian National Research Council (CNR). The selected proposal "ECHOES - European Cloud for Heritage OpEn Science" sets out the challenge of building a platform common to all other subsequent projects (currently more than 10) that have been and will be launched by the EC with other calls.



With a broad multidisciplinary partnership (51 partners), ECHOES will design and implement the basic tools for building and managing digital applications for the Cultural Heritage. This digital ecosystem of data, components and tools will enable multiple consortia to develop and share many applications tailored to the needs of the different communities of CH professionals. Unlike in previous attempts, the resources will be both effective and practical, yet able to be part of a common rich digital ecosystem.

The common ecosystem will support: interactive, cross-platform visualization of visual and textual data; joint design of common user interfaces; instruments for evaluating the effectiveness and usage of applications; and, finally, a few sample applications that can also serve as examples for other consortia winning any subsequent ECCCH call.

These efforts will be complemented by the usual training and dissemination activities aimed at stakeholders.

A final task of ECHOES will be to create a legal entity and a policy for the future management of the cooperative cloud infrastructure to ensure its permanence.

Italy's participation in ECHOES includes the CNR, with the institutes ISTI (Coordinator), ISPC, ILC, and ICMATE, the research centers CRS4 (Cagliari) and IIT (Genoa), involved as subcontractors for CNR, PIN (Prato) and the Italian ETT/META Group.

Contact: Roberto Scopigno, Director
roberto.scopigno@isti.cnr.it

Generative AI for Images and Virtual Worlds

This article overviews a key scientific and technological advancement with huge societal impact: Generative Artificial Intelligence for creating images, videos, and 3D content such as virtual worlds. We briefly review the existing systems and discuss major ethical issues, which in turn inspire novel and interesting directions for future scientific research.

Introduction

Unlike traditional machine learning methodologies, systems based on Generative AI (GenAI) aim to do more than analyse or categorize data. They strive to capture the underlying patterns, structures, and relationships defining a specific dataset.

The neural networks constituting these systems consist of layers that work together to extract and refine data. As information passes through the layers, the system identifies increasingly complex and abstract features, gradually building a deeper “understanding” of the dataset with each round of training.

The key goal of these systems is to model the probability distribution of the data they are trained on. When exposed to large datasets — be these images, text, or multimedia — they create a statistical model capturing the traits, variations, and patterns within the data. This enables them to generate new outputs reflecting the same features.

Societal and scientific impact

The impact of generative AI is already evident in real-world applications across science and the arts, offering tools that complement and improve human expertise.

In medicine, for instance, researchers are using AI to analyse complex molecular structures, simulate drug interactions, and propose new compounds. These steps would have taken far longer without such technologies. Similarly, in environmental science, AI aids in understanding climate change by

creating detailed models and running simulations, providing insights that were previously unattainable. In material science, AI facilitates the design of innovative materials with the potential to transform areas such as clean energy, advanced electronics, and environmental sustainability. Although these breakthroughs still require rigorous scientific validation, they significantly accelerate the discovery process.

Generative AI is also pushing the boundaries of creativity. It is being used to create music, stories, illustrations, and stunning photorealistic images. These tools are not intended to replace human imagination but to provide new ways to explore ideas and expand what is possible.

A critical innovation underpinning state-of-the-art systems for image and video generation (e.g., Stable Diffusion, Flux, Midjourney, Dall-E, and Sora) is the notion of a diffusion model. This model is enhanced by the so-called latent diffusion process, which operates on a compressed representation of the data. This makes the problem computationally feasible on consumer-grade hardware.

Diffusion models

Diffusion models represent a groundbreaking advancement in generative AI. They offer a novel approach to creating high-quality synthetic data, including images, videos, and even 3D content. The foundation of these models lies in their ability to capture underlying patterns and structures within a dataset by simulating a reverse diffusion process. This process involves gradually denoising a random noise distribution to reconstruct the target data, enabling the generation of remarkably realistic outputs.

The training process of diffusion models begins with progressively adding noise to the original data, effectively corrupting it over multiple steps. The model learns to reverse this process step by step, gradually refining noisy inputs into coherent outputs. By iteratively

minimizing the difference between the generated and real data across these steps, diffusion models achieve a highly accurate representation of the underlying data distribution. This iterative framework, inspired by concepts in statistical thermodynamics, has proven highly effective for generating data with fine-grained details and variations. One key innovation of diffusion models is their ability to operate in a latent space rather than directly in the pixel or feature space of the original data. By compressing high-dimensional data into a lower-dimensional latent representation, diffusion models achieve both computational efficiency and improved generation quality. This approach allows models like Stable Diffusion to synthesize complex images while maintaining a manageable resource footprint, making the technology more accessible for broader applications.

Diffusion models have also been adapted for conditional generation, where additional information — such as text prompts, class labels, or reference images — guides the generation process. This capability has unlocked new possibilities for creative applications. These include generating personalized artwork, realistic portraits, and virtual worlds with contextually consistent elements. The adaptability of diffusion models has made them a cornerstone in the evolving landscape of generative AI technologies.

Current systems

Open source: Stable Diffusion and Flux

Stable Diffusion is a state-of-the-art model designed for generating highly realistic images from textual descriptions. Built upon the principles of latent diffusion, it represents a significant advancement in image synthesis technology. Stable Diffusion was primarily developed by Stability AI, with contributions from academic and industrial collaborators, including the CompVis group at the Ludwig Maximilian University of Munich (formerly at the University of Heidelberg).

One of the key advantages of Stable Diffusion is its accessibility. Unlike earlier models that required extensive computational resources and specialized hardware, it is optimized to run on consumer-grade hardware, bringing powerful image-generation capabilities to a broader audience. This democratization of advanced AI technology has opened up numerous applications, including art creation, design, content generation, video synthesis, and more. The current version of Stable Diffusion is 3.5, featuring high-quality images, and strong prompt adherence.

Recently, a class of models published by Black Forest Labs, called FLUX.1, has set a new benchmark in image synthesis. These models are based on a multimodal architecture called diffusion transformers. The “pro” version (available only via API and servers) outperforms leading open and closed-source models in key areas such as visual quality, prompt adherence, support for diverse sizes and aspect ratios, typography handling, and output diversity. The “dev” version is derived from the “pro” version via so-called “distillation,” and its weights are free to download for non-commercial use. The “schnell” variant, which is open source, stands out as the most advanced existing model available that runs in a few steps, surpassing other models in its category and rivaling the quality of higher-performing non-distilled options.

Proprietary: Dall-E, Midjourney etc

Most closed-source systems (all of which are web applications living in the cloud) prioritize usability and simplicity, making them highly appealing to a broad user base. These platforms enable users to generate high-quality visuals by providing simple textual prompts, ensuring accessibility even to those without technical expertise in AI or image generation.

However, the proprietary nature of these systems introduces certain limitations. Transparency is often restricted, and the creative possibilities are reduced by strict

prompt filtering. This filtering aims to avoid generating questionable and unethical material but can also prevent the creation of genuine content requiring specific, filtered keywords (e.g., the representation of nudity or violence in an artistic context).

The underlying models, data sources, and training methodologies are typically not disclosed. This opacity can lead to challenges in understanding how the systems work and addressing biases or inaccuracies in the generated content. Additionally, these platforms often require payment, either through subscription fees or usage-based pricing, which can limit accessibility for individuals, small organizations, and educational institutions.

Two key tools are i) Midjourney (<https://midjourney.com/>), whose extremely high quality of its images, make it popular with artists, designers, and creative professionals, and ii) Dall-E (<https://openai.com/dall-e-2/>) which is simple to use and is integrated with ChatGPT.

The language(s) of Generative AI

An interesting theme for future research is that AIs can be composed and programmed to operate together. One might argue that this is a temporary phenomenon, as more complex systems will eventually be entirely programmed via textual prompts. However, both open and closed-source models will continue to be produced by a multitude of different entities.

Just as humans coordinate efforts through protocols and guidelines – which may be more or less formalized depending on the context – AI systems also need to be orchestrated, composed, or programmed at scale. A key point to remember is that the users of these systems are not, and likely will not be in the foreseeable future, traditional computer programmers. These users will require tools that empower them with programming and testing/monitoring abilities, enabled through novel programming language paradigms.

The last three or four decades of research in the fields of logics and formal methods in computer science have produced extensive literature devoted to service composition, modern type theories encompassing dynamic computation aspects, and specification languages that express system properties without requiring direct programming.

From the perspective of human-computer interaction, advancements in programming editors have made rapid prototyping and system deployment more accessible to many programming-literate users. It is time to redirect this impressive body of knowledge toward Generative AI, designing novel, visual systems that are simultaneously easy to use and program. Through modern user interfaces, programming can become nearly invisible while retaining the benefits of theoretical rigor. For instance, safe and transparent coordination of different closed-source services can be achieved. Moreover, type systems could limit errors, simplifying programming by narrowing choices. Declarative specification languages could be employed quickly verify whether produced data adheres to quality constraints.

One such system is being developed at CNR-ISTI. The VoxLogicA project (www.voxlogica.org) aims to define a domain-specific, executable meta-language (https://link.springer.com/chapter/10.1007/978-3-031-75387-9_13), that facilitates the composition of diverse services operating in imaging and related datasets. This meta-language encompasses constructs central to traditional AI, such as spatial and temporal logics, alongside typical machine learning operations like training, prediction, fine-tuning, and dataset augmentation.

The project also emphasizes human-computer interaction from the start (https://link.springer.com/chapter/10.1007/978-3-031-06388-6_14) designing graphical user interfaces driven by principles of cognitive load minimization and dataset orientation.

Generative AI for Virtual worlds

Introduction to Virtual Worlds in the Context of Generative AI

Generative AI has revolutionized how we create, understand, and interact with virtual worlds, particularly through the lens of 3D assets. These assets – digital representations of objects, environments, and entities in three-dimensional space – serve as the foundation for two critical applications: generating robust training datasets for AI models and enabling context-aware, statistically plausible 3D world completion. Together, these advancements are shaping new frontiers in extended reality (XR) exploration, where immersive and adaptive environments are increasingly vital.

Virtual worlds augmented by generative AI and 3D assets are redefining how we interact with digital spaces. By powering training dataset creation and enabling dynamic world completion, these technologies open the door to more realistic, responsive, and adaptive experiences.

3D Assets for AI Training Dataset Generation

In generative AI, the creation of high-quality training datasets is essential for developing models capable of understanding and replicating complex spatial dynamics. Virtual worlds populated with synthetic 3D assets provide an efficient, scalable solution. These assets allow for the simulation of diverse scenarios that would be difficult, expensive, or impossible to replicate in the real world. For example, AI models can be trained on various environmental conditions, object interactions, or crowd dynamics within these synthetic spaces. The controlled variability of virtual worlds ensures datasets are rich and balanced, significantly improving model robustness and generalization to real-world applications.

Generative 3D Completion for Partially Known Worlds

In XR exploration, users often encounter



Images generated using the Stable Diffusion 1.5 workflow in ComfyUI (see <https://github.com/comfyanonymous/ComfyUI>) and the Stable Diffusion 1.5 fine-tuned model DreamShaper 8 (see <https://huggingface.co/Lykon/DreamShaper>). Images have been upscaled using the Real-ESRGAN 8x upscaler (model available at <https://huggingface.co/ai-forever/Real-ESRGAN>). Prompts: 1) "Einstein taking a selfie using a smartphone, standing, full figure, black and white photo, midshot" 2) "sad kangaroo looking straight into the camera, wearing a hoodie, portrait, pulitzer, urban, closeup, lens" 3) "Artificial intelligence, robot, abstract:1.2, bionic, black woman, computer digital image, profile, closeup profile, glass skin, cyberpunk, implants, digital circuit brain, cyborg" 4) "Michelangelo's David playing a (guitar made of marble):1.5, sitting, thoughtful, marble sculpture".

scenarios where the available 3D data is incomplete or partially obscured. Examples include virtual environments where only parts of a building or terrain are mapped. Generative AI addresses this challenge by leveraging probabilistic models and learned spatial patterns to infer the missing details.

For instance, neural networks trained on large datasets of 3D assets can predict the likely structure of an incomplete object or environment. This generates realistic, context-aware completions that seamlessly integrate with the known elements of the world. This is especially true when representing space with image as a volume data structures, especially by encoding the interface between the object and the exterior as a signed distance field, as shown in <https://iris.cnr.it/handle/20.500.14243/493141> and in the recent <https://iris.cnr.it/handle/20.500.14243/514771>.

Applications and Implications

The intersection of 3D asset generation and generative AI has profound implications across industries. In architecture and urban

planning, these technologies enable detailed simulations of yet-to-be-built structures in complex urban settings. In gaming and entertainment, they enhance procedural generation, creating expansive and interactive virtual worlds. For XR applications, such as virtual tourism or remote training, generative AI transforms partially mapped environments into immersive, navigable experiences.

Moreover, the statistical plausibility offered by generative models ensures reconstructed or synthesized elements adhere to real-world physics, aesthetics, and functional patterns. This capability fosters a sense of realism and trust, critical for applications where user engagement and accuracy are paramount.

Ethical issues, and future directions

Generative AI technologies for images and videos come with significant ethical considerations. These issues can be grouped into three primary categories, each representing a direction for novel and important research.

Deep fakes

Deep fakes are AI-generated images or videos that convincingly depict people or events that never occurred. These media are often created by synthesizing or manipulating content to mimic real individuals.

The primary issues with deep fakes include the spread of false information, the erosion of trust, harassment (especially through non-consensual explicit content), political or financial fraud, and identity theft. Despite the rapid development of generative AI, laws to regulate its misuse often lag behind. Governments are drafting regulations to penalize malicious creation and distribution of deep fakes.

In the interim, strategies like watermarking (embedding detectable signatures into AI-generated content) and deep learning-based detection tools can mitigate some risks. However, this sets up an ongoing race between those misusing AI and those developing countermeasures.

Further reading - Active research is being conducted at CNR-ISTI on the matter of deepfakes. See e.g. <https://peerj.com/articles/cs-2127/> and the citations therein. The following RANDR report is a good starting point to study these matters: <https://www.rand.org/pubs/perspectives/PEA1043-1.html>.

Copyright issues

Generative AI systems are frequently trained on large datasets scraped from the internet, which can include copyrighted materials such as images, videos, and text. These systems sometimes inadvertently or intentionally replicate copyrighted content. Key concerns involve fair use violations, plagiarism, and the economic impact on human-generated art and design, potentially devaluing original works.

One proposed mitigation strategy involves transparent training data. Ethical guidelines should be developed to ensure datasets respect intellectual property.

Central to this approach are opt-out mechanisms and compensation models for creators whose work is used.

Further reading - See, e.g., MIT Technology Review's article on opt-out mechanisms <https://www.technologyreview.com/2022/12/16/1065247/artists-can-now-opt-out-of-the-next-version-of-stable-diffusion/>.

Pornography and child abuse.

A pressing ethical concern in open-source generative AI for images and videos is its potential misuse in creating pornographic material. Alarming, child sexual abuse material (CSAM) has been inadvertently included in some public datasets, such as LAION-5b, used to train diffusion models.

While this particular case is largely considered an unintentional oversight, it highlights the urgent need for robust measures. Perpetrators exploit these technologies, sometimes falsely claiming real images are AI-generated. Law enforcement faces increasing challenges in identifying and addressing such cases as AI-generated content proliferates. Furthermore, the task of locating and aiding actual victims becomes substantially more challenging, since efforts are diverted towards generated material. Advanced classifiers and better regulation are necessary to curb this misuse.

Further reading - https://www.iwf.org.uk/media/drufozvi/iwf-ai-csam-report_update-public-jul24v12.pdf, <https://arstechnica.com/tech-policy/2024/08/nonprofit-scrubs-illegal-content-from-controversial-ai-training-dataset/>.

Conclusions

While diffusion models and other generative AI technologies hold immense potential for innovation, they also pose significant ethical challenges. Addressing these issues requires a multifaceted approach involving technological safeguards, ethical guidelines, legal reforms, and public awareness. Collaboration among researchers, developers, pol-

icymakers, and society is essential to ensure these technologies are used responsibly and for the betterment of humanity.

Generative AI has “democratized creativity,” enabling anyone to produce high-quality images, music, and videos that rival professional outputs. However, like many disruptive technologies, its adoption may have unintended economic consequences for artists, designers, and content creators. Judging whether the positive effects outweigh the negatives remains complex.

Despite these challenges, the open problems and applications described make generative AI an exciting research field and a valuable support tool for neighbouring disciplines like image analysis. We hope this article inspires ISTI researchers to explore new directions and projects in this domain.

Contacts: Vincenzo Ciancia, FMT Lab
vincenzo.ciancia@isti.cnr.it

Marco Di Benedetto, AIMH Lab
marco.dibenedetto@isti.cnr.it

Carlo Metta, KDD Lab
carlo.metta@isti.cnr.it

MMinE-SwEEPER

Marine Munition in Europe - Solutions with Economic and Ecological Profits for Efficient Remediation Funded by Horizon Europe

Explosive and chemical munitions dumped in the sea pose a threat to marine traffic, infrastructure, humans, and the environment. Clearing unexploded ordnance (UXO) is essential for the safety and security of European waters. The EU-funded MMinE-SwEEPER project aims to enhance Europe's ability to manage marine munitions in non-military scenarios. The project will focus on advancing technologies for automated munition detection and identification, data analysis, environmental monitoring of chemical contaminants, predicting UXO burial and the spread of contaminants, assessing corrosion, and enabling the secure exchange of sensitive data. It will also involve stakeholder dialogue and outreach activities. Industry partners anticipate significant benefits from these technological developments, and governmental bodies are looking forward to closer cooperation between European countries.

The MMinE-SwEEPER project will advance the knowledge, skills and capacity in Eu-

rope for dealing with marine munitions in the non-military aspects of UXO-clearance. A total of 21 partners from 7 EU and 2 EU-associated countries bring experience and skills from the civil science community, military research, coast/border guards or EOD-services as well as industry and intergovernmental organizations. These partners will work together in 8 technical work packages to a) advance automated munition detection, identification and data analyzing technologies and software, b) provide environmental monitoring of chemical contaminants, c) predict UXO-burials, the spread of contaminants and state-of-corrosion, and d) enable the secure exchange of sensitive data. Three additional WPs will gather the existing knowledge, learn about the legal responsibilities and technical approaches, assess remediation and mitigation measures, and create training materials in order to build the European capacity. Two WPs will engage in an meetings with stakeholders and dissemination/outreach by using HELCOM and JPIO resources

Outcomes with respect to specific technical advancements will include a) AI-supported detection of munition in hydroacoustic spatial mapping data (MBES, SSS; SAS), b) detection of buried munition objects (magnetic, SBP, LF-SAS), c) AI-supported object identification in optical and acoustic cameras and d) incorporating trained AI-models within Smart-AUVs and -USVs for adaptive and cooperative mission execution. Secure data exchange will be provided through a demonstrator data platform that will be refined during the project. Industry partners see great benefits in developing technologies that can be further advanced to become commercial products. At the end of the project most TRLs will score 5 to 6, which means that they have been applied under real conditions and proved their applicability.

Contact: Massimiliano Corsini, VC Lab
massimiliano.corsini@isti.cnr.it



MUSIT

Multi-Sensor Inferred Trajectories Funded by Horizon Europe

The abundance of tracking sensors in recent years has led to the generation of high-frequency and high-volume data streams, including vessels, vehicle tracking data, smart-watches, cameras, and earth observation sensors. However, the trajectory of a moving object may not be complete, contain errors, or simply not be unavailable. Although a vast pool of tracking data is available, it can remain unexplored or underutilized, and it could reveal important information.

The Multi-Sensor Inferred Trajectories (MUSIT) project merges data from heterogeneous sources in order to provide detailed information on a moving object's whereabouts and behavior, reducing gaps, and producing a refined and inferred trajectory with minimal errors. The fusion of multi-sensor data is required to fill in the trajec-

tory gaps of moving objects and to attach useful semantics to the trajectory and its components.

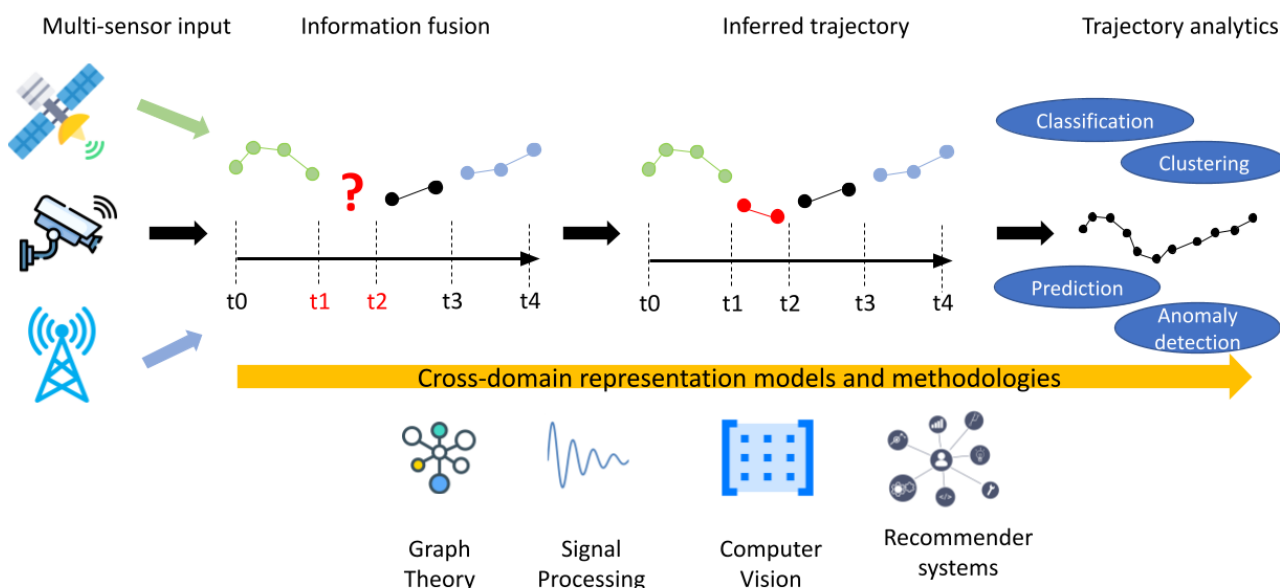
In implementing MUSIT, AI algorithms and spatio-temporal methodologies that can fuse information and infer the “missing knowledge” are crucial. In addition, different representation models from multiple domains within the ICT sector will also be explored. Datasets will be made available in cases that were previously thought to be impossible and will infer knowledge, thus improving the overall surveillance.

MUSIT project will be divided into three parts: i) data collection and creation, ii) exploitation and use of cross-domain representation models within the ICT sector for trajectories, and iii) analysis and processing

of outcomes to produce information-rich results related to vessel monitoring and urban mobility.

Results from MUSIT will be reported in open access journals. In fact, MUSIT will be part of the open science community, promoting the adoption and support of FAIR principles in the fields of AI and data analysis. In addition it advocates guiding principles for good research data management, which will lead to a higher degree of reusability and sustainability of data and other research artifacts.

Contact: Emanuele Carlini, HPC Lab
emanuele.carlini@isti.cnr.it
<https://www.musit-project.eu/>



Indoor Localization Algorithms Based on Angle of Arrival with a Benchmark Comparison

F. Furfari, M. Girolami, F. Mavilia, P. Barsocchi
Ad Hoc Networks, vol. 166. Elsevier, 2024.

Indoor localization is crucial for developing intelligent environments capable of understanding user contexts and adapting to environmental changes. Bluetooth 5.1 Direction Finding is a recent specification that leverages the angle of departure (AoD) and angle of arrival (AoA) of radio signals to locate objects or people indoors. This paper presents a set of algorithms that estimate user posi-

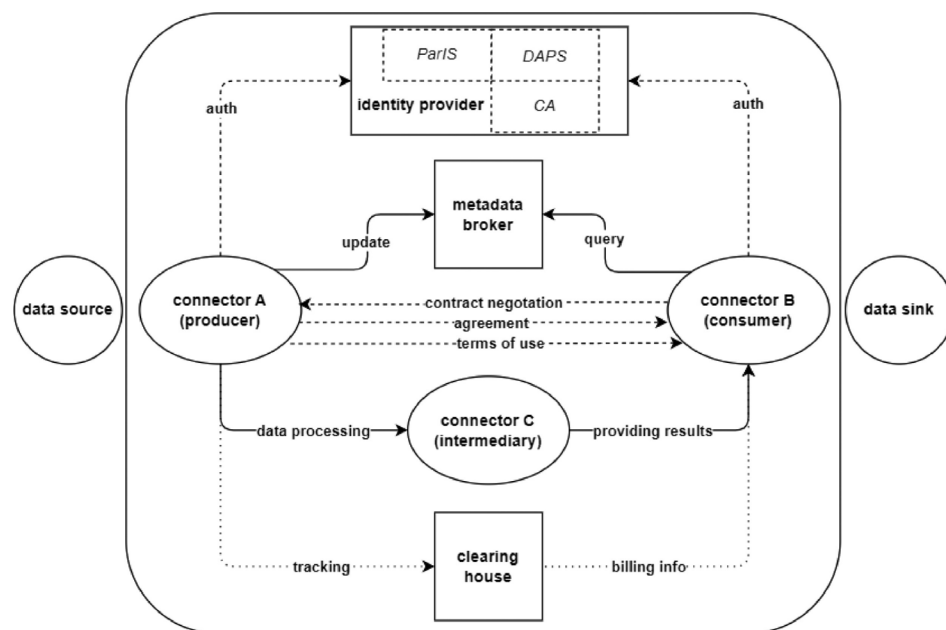
tions using AoA values and the concept of the Confidence Region (CR), which defines the expected position uncertainty and helps to remove outlier measurements, thereby improving performance compared to traditional triangulation algorithms. We validate the algorithms with a publicly available dataset, and analyze the impact of body orientation relative to receiving units. The experi-

mental results highlight the limitations and potential of the proposed solutions. From our experiments, we observe that the Conditional All-in algorithm presented in this work, achieves the best performance across all configuration settings in both line-of-sight and non-line-of-sight conditions.

DOI: 10.1016/j.adhoc.2024.103691

What are Data Spaces? Systematic Survey and Future Outlook

M. Bacco, A. Kocian, S. Chessa, A. Crivello, P. Barsocchi
Data in Brief, vol. 57. Elsevier, 2024.



Interactions among the main technical components in a data space. Dotted lines are for control and administrative messages, and solid lines are for data and metadata exchanges.

Data spaces, a novel concept pushing data sharing and exchange, are experiencing momentum because of recent develop-

ments motivated by the increasing need for interoperability and data sovereignty. After an initial phase, dating back to ap-

proximately twenty years ago, in which this concept has been tentatively explored in different scenarios, it is presently going through a consolidation phase in which both specifications and implementations converge towards a common reference for standardisation. In this context, we offer our view on data spaces by presenting a systematic literature survey, a description of the components needed to build them, how they work, and of existing mature software implementations. We thoroughly present the architectural vision behind the concept and we analyse the Reference Architectural Model by IDS. We provide practical pointers to readers interested in experimenting with software components used in data spaces, and we conclude by highlighting open challenges for their success.

DOI: 10.1016/j.dib.2024.110969

Formal Methods in Industry

M.H. ter Beek, R. Chapman, R. Cleaveland, H. Garavel, R. Gu, I. ter Horst, J.J.A. Keiren, T. Lecomte, M. Leuschel, K.Y. Rozier, A. Sampaio, C. Seceleanu, M. Thomas, T.A.C. Willemse, L. Zhang
Formal Aspects of Computing, vol. 37. ACM, 2024.

Formal methods encompass a wide choice of techniques and tools for the specification, development, analysis, and verification of software and hardware systems. Formal methods are widely applied in industry, in activities ranging from the elicitation of requirements and the early design phases all the way to the deployment, configuration, and runtime monitoring of actual systems. Formal methods allow one to precisely specify the environment in which a system operates, the requirements and properties that the system should satisfy, the models of the system used during the various design steps, and the code embedded in the final implementation, as well as to express conformance relations between these specifications. The key motivation for these techniques is the strength of the correctness guarantees they provide: in contrast to testing-based and inspection-based techniques, a proof conclusively demonstrates that the system in question, at the level of abstraction that it is presented, is correct with respect to its specification. Formal methods complement other verification and validation techniques, such as testing or simulation.

This paper demonstrates that formal methods have wide-ranging practical value by reporting on the increasing use of formal methods in industry, and it makes a case for the inclusion of formal methods as a separate topic in Computer Science education. This is well agreed upon by the formal methods community. To make this point better known to the Computer Science community at large and, in particular, to those involved in Computer Science education, we present a broad scope of successful applications of formal methods in industry, not limited to

the well-known success stories from the safety-critical domain, like railways and other transportation systems, but also covering other areas such as lithography manufacturing and cloud security in e-commerce, to name but a few. We also report testimonies from a number of representatives from industry who, either directly or indirectly, use or have used formal methods in their industrial project endeavours. These persons are spread geographically, including Europe, Asia, North and South America, and the involved projects witness the large coverage of applications of formal methods, not limited to the safety-critical domain.

In this paper, we thus make a case for the importance of formal methods, and in particular of the capacity to abstract and mathematical reasoning that are taught as part of any formal methods course. We argue that undergraduate curricula should include formal methods as a topic, not only because of their growing importance in industry, as witnessed by the evidence presented in the paper, but also because formal methods contribute to cultivate abstract thinking, enabling students to better understand and solve complex problems, and because of the discipline they instil in students as they learn to develop systems. These are fundamental Computer Science skills that graduates should profit from when working as computer scientists in industry, as confirmed by several industry representatives involved in the writing of this paper. In the words of Rance Cleaveland, who died during the revision of this paper and whom we honour with this paper, “students who learn formal methods are better developers, because they learn to think about correctness while they are building systems.”

Further support was recently received from The White House Office of the National Cyber Director in the technical report “Back to the Building Blocks: A Path Toward Secure and Measurable Software” (<https://www.whitehouse.gov/wp-content/uploads/2024/02/Final-ONCD-Technical-Report.pdf>), which advocates the use of formal methods over testing for demonstrating the correctness of software and considers it vital to make formal methods widely accessible to accelerate broad adoption (cf. Part II: Securing the Building Blocks of Cyberspace—Formal Methods): “Given the complexities of code, testing is a necessary but insufficient step in the development process to fully reduce vulnerabilities at scale. If correctness is defined as the ability of a piece of software to meet a specific security requirement, then it is possible to demonstrate correctness using mathematical techniques called formal methods. These techniques, often used to prove a range of software outcomes, can also be used in a cybersecurity context and are viable even in complex environments like space. While formal methods have been studied for decades, their deployment remains limited; further innovation in approaches to make formal methods widely accessible is vital to accelerate broad adoption. Doing so enables formal methods to serve as another powerful tool to give software developers greater assurance that entire classes of vulnerabilities, even beyond memory safety bugs, are absent.”

DOI: 10.1145/3689374


The Role of Formal Methods in Computer Science Education

M.H. ter Beek, M. Broy, B. Dongol.
ACM Inroads, vol. 15. ACM, 2024.

In this article, we summarise discussion and evidence brought forward in three white papers—all together authored by 35 computer scientists and practitioners worldwide—on the key position of Formal Methods in Computer Science (CS) education.

Formal Methods have multiple characterisations in the literature as languages and techniques (and tools) based on rigorous mathematical foundations for the specification, development, and (manual or automated) analysis and verification of software and hardware systems. According to IEEE, “software engineering methods provide an organized, systematic approach for specifying, designing, constructing, testing, and verifying the resulting software products and associated work items involved in developing computer software applications. The methods impose a certain structure, set of steps, practices, and procedures on the software engineering effort to make it more methodical, repeatable, and more success-oriented” (<https://www.computer.org/resources/software-engineering-models>). A method is called “formal” if its set of steps is applied using formal techniques such that the correctness of the result of the application is formally justified.

The 35 authors of the aforementioned white papers have extensive experience in teaching Formal Methods in Academia or working with them in Industry. They argue for introducing Formal Methods as a distinct core knowledge area in the ACM/IEEE/AAAI CS curricula guidelines (<https://csed.acm.org/>),



Teaching Formal Methods (Image courtesy of Luigia Petre, generated using Dall-E).

since currently Formal Methods are covered only as elective topics in some knowledge units of the knowledge areas Foundations of Programming Languages and Software Engineering. The evidence put forward in these papers suggests that the lack of indicating Formal Methods as a core knowledge area in the CS2023 guidelines is incomprehensible, arguably similar to the fact that Artificial Intelligence was not included just 10 years ago in the CS2013 guidelines.

First, ter Beek et al. provide an extensive review of the increasing use of Formal Methods in Industry—not limited to the well-known success stories from the safety-critical domain. For example, they present use cases in lithography manufacturing and cloud security in e-commerce, and industrial use of Formal Methods in Europe, Asia, and North and South America. Second, Dongol

et al. stipulate the importance of Formal Methods thinking in CS education, since this provides the necessary rigour in reasoning about software, its specification, its verification, and its correctness—all fundamental skills for future software developers. Third, Broy et al. advocate the importance, for all computer scientists, of knowing Formal Methods, since the skills and knowledge acquired in this way provide the indispensable solid foundation that forms the backbone of CS practice.


Finally, we underline that teaching Formal Methods need not come at the cost of displacing other engineering aspects of CS that are already widely accepted as essential. On the contrary, we show that Formal Methods have the potential to support and strengthen the presentation and knowledge in eight of the 17 knowledge areas of CS2023; next to the abovementioned knowledge areas, these are Algorithmic Foundations, Architecture and Organization, Artificial Intelligence, Parallel and Distributed Computing, Security, and Software Development Fundamentals. We moreover provide concrete suggestions for educators on how to incorporate Formal Methods into CS education, for example via Formal Methods Europe (FME, <https://fmeurope.org/>) and its Teaching Committee’s Education Course Database (<https://fme-teaching.github.io/courses/>).

DOI: 10.1145/3702231

Advancing Orchestration Synthesis for Contract Automata

D. Basile, M.H. ter Beek

Journal of Logical and Algebraic Methods in Programming, vol. 141. Elsevier, 2024.



Railway scenario, with one train waiting to enter a junction area while another train is traversing it.

Contract automata allow to formally define the behaviour of service contracts in terms of service offers and requests, of which some are moreover optional and some are necessary. A composition of contracts is said to be in agreement if all service requests are matched by corresponding offers. Whenever a composition of contracts is not in agreement, it can be refined to reach an agreement using the orchestration synthesis algorithm. This algorithm is a variant of the synthesis algorithm used in supervisory control theory and it is based on the fact that optional transitions are controllable (i.e., such transitions can be blocked by the controller), whereas necessary transitions are at most semi-controllable and cannot always be controlled.

In this paper, we present advancements of the orchestration synthesis for contract automata. Notably, we identify the existing limits of the orchestration synthesis and propose a novel orchestration synthesis along with additional constructs to enhance the expressiveness and scalability of contract automata. The proposed advancements have been implemented and experimented on two case studies, one of which is a card game while the other is based on a real-world scenario from the railway domain.

The railway case study concerns an autonomous train positioning (ATP) system, in which the physical track circuits detecting the occupancy of portions of the railway

track are substituted by virtual track circuits (VTCs). These VTCs are virtual positions on a map. The real position of a train is detected using a global positioning system (GPS). In the specific scenario, depicted in Fig. 1, one junction area (commanded by one interlocking) is composed of two VTCs, with one train outside and one train inside the junction area. Train 2 is traversing its assigned route, while Train 1 is waiting at a red signal for its route to be assigned. VTC 0 is used to detect the occupation of a route, whereas VTC 1 is used to detect the release of a route. Initially, both trains are located behind the semaphore. The first train arriving (Train 2) will communicate its route to the interlocking, which will proceed to set the route. This may cause the movement of the junction point. Once the route is set, the interlocking will signal to Train 2 that the route is set by opening the semaphore. Train 2 enters the junction point and the semaphore is closed again. While Train 2 is traversing its route, the second train arriving (Train 1) will stop at the (closed) semaphore to ask for its route. The route will be assigned, the junction point moved, and the semaphore opened only after Train 2 has exited the junction area. Otherwise, the movement of the junction point could cause the derailment of Train 2 inside the junction area.

In this paper, we show how to model this scenario with five contract automata. Two train automata are such that their states represent spatial coordinates in a bi-dimen-

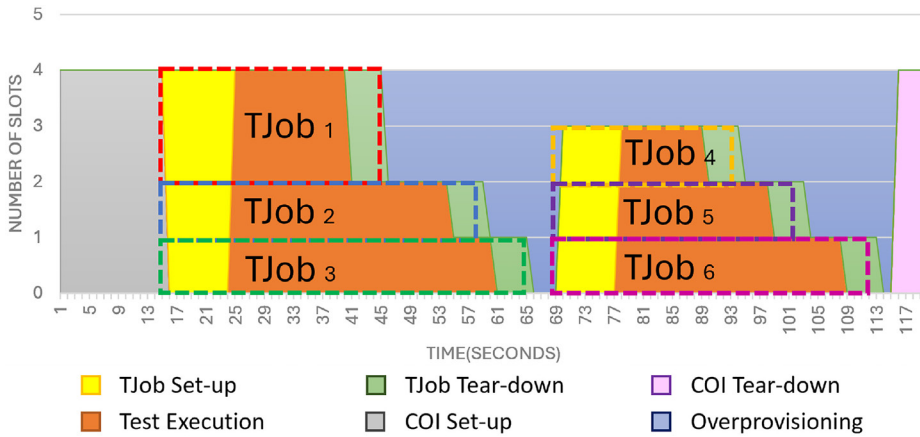
sional map (the railway track depicted in Fig. 1). The trains request to move to an adjacent location through necessary lazy requests (i.e., a train is semi-controllable). The junction area is identified as a specific location, where a train is allowed to perform a specific transition modelling the traversing of the point. A driver automaton is used to match the requests of the trains to move. Similarly, two automata control the semaphore by opening and closing it. When composing these five contract automata, we enforce agreement and identify the forbidden states in the composition. The forbidden states are those that satisfy one or more of the following requirements: (i) both trains are in the same location (and not in the final location modelling the exit from the junction area); (ii) one train is at the semaphore location while the semaphore is closed; (iii) both trains are inside the junction area; (iv) a train is inside the junction area and the semaphore is opened; and (v) the semaphore is opened but no train is near it (in this case the semaphore must be closed). These requirements are invariants that can be checked on a state of the composition during state space generation. The orchestration of this case study is a strategy for the semaphore controller to command the opening and closing of the semaphore in such a way that both (semi-controllable) trains are allowed to reach the exit while satisfying the above requirements. For the railway case study, the novel orchestration strategy compares favourably for what concerns performance and size.

DOI: 10.1016/j.jlamp.2024.100998

This work was carried out within MUR PRIN 2022 PNRR P2022A492B project ADVENTURE (ADVancEd iNtegrated evalUation of Railway systEMs) and the MOST - Sustainable Mobility National Research Center and received funding from the European Union NextGenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) - MISSIONE 4, COMPONENTE 2, INVESTIMENTO 1.4 - D.D. 1033 17/06/2022, CN00000023).

RETORCH*: A Cost and Resource Aware Model for E2E Testing in the Cloud

C. Augusto, J. Morán, A. Bertolino, C. de la Riva, J. Tuya
Journal of Systems and Software, vol. 221. Elsevier, 2024.



Example of an as-a-service web browsers Usage Profile.

Moving testing to the Cloud overcomes time/resource constraints by leveraging an unlimited and elastic infrastructure, especially for testing levels like End-to-End (E2E) that require a high number of resources and/or execution time. However, it intro-

duces new challenges to those already faced on-premises, like selecting the most suitable Cloud infrastructure and billing scheme. We propose the RETORCH* test execution model that estimates and compares the monetary cost of executing an E2E test

suite with different Cloud alternatives, billing schemes, and test configurations. RETORCH* goes beyond the mere cost billed, and selects the solution that best aligns with the test team strategy using the data of on-premises prior executions and the tester's experience. This cost is broken down into the cost incurred to execute the test suite (testing cost) and possible unused infrastructure (overprovisioning cost). Based on these distinct costs, the test team can compare different Cloud and test configurations. RETORCH* has been evaluated using a real-world application's E2E test suite. We analyze how the different decisions taken when the suite is migrated to the Cloud impact the cost, highlighting how RETORCH* can help the tester during Cloud and test configuration to make a more informed decision.

DOI: 10.1016/j.jss.2024.1122371

A Risk Assessment Framework Based on Fuzzy Logic for Automotive Systems

F. Merola, C. Bernardeschi, G. Lami
Safety, vol. 10. MDPI, 2024.

Recent advancements in the automotive field have significantly increased the level of complexity and connectivity of modern vehicles. In this context, the topic of cybersecurity becomes extremely relevant, as a successful attack can have an impact in terms of safety on the car navigation, potentially leading to harmful behavior. Risk assessment is typically performed using discrete input and output scales, which can often lead to an identical output in terms of risk evaluation despite the inputs present-

ing non-negligible differences. This work presents a novel fuzzy-logic-based methodology to assess cybersecurity risks which takes attack feasibility and safety impact as input factors. This technique allows us explicitly model the uncertainty and ambiguity of input data, which is typical of the risk assessment process, providing an output on a more detailed scale. The fuzzy inference engine is based on a set of control rules expressed in natural language, which is crucial to maintaining the interpretability

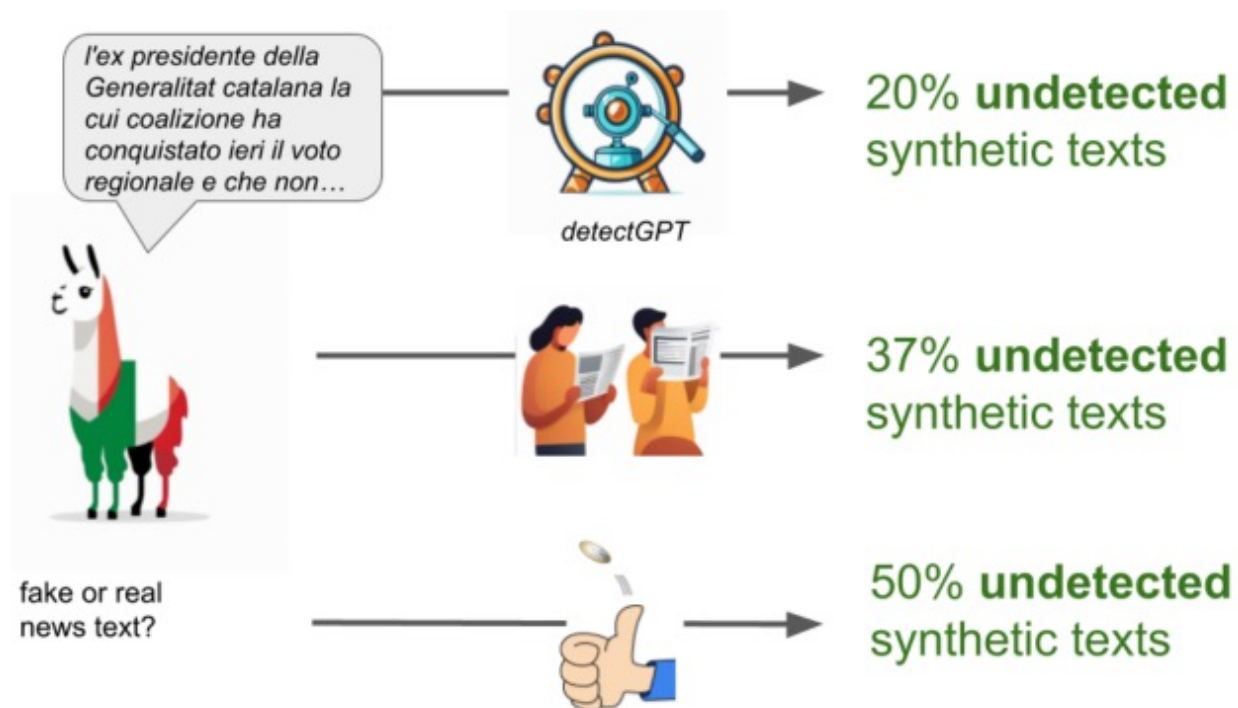
and traceability of the risk calculation. The proposed framework was applied to a case study extracted from ISO/SAE 21434. The obtained results are in line with the traditional methodology, with the added benefit of also providing the scatter around the calculated value, indicating the risk trend. The proposed method is general and can be applied in the industry independently from the specific case study.

DOI: 10.3390/safety10020041

AI ‘News’ Content Farms are Easy to Make and Hard to Detect: a Case Study in Italian

G. Puccetti, A. Rogers, C. Alzetta, F. Dell'Orletta, A. Esuli

Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics. ACL, 2024.



Detecting synthetic Italian news text generated by fine-tuned Llama-65B: error rates for DetectGPT, native speakers of Italian, and random guess.

Large Language Models (LLMs) are increasingly used as ‘content farm’ models (CFMs), to generate synthetic text that could pass for real news articles. This is already happening even for languages that do not have high-quality monolingual LLMs. We show that fine-tuning Llama (v1), mostly trained on English, on as little as 40K Italian news articles, is sufficient for producing news-like texts that native speakers of Italian struggle to identify as synthetic. We investigate three LLMs and three methods of detecting synthetic texts (log-likelihood, DetectGPT,

and supervised classification), finding that they all perform better than human raters, but they are all impractical in the real world (requiring either access to token likelihood information or a large dataset of CFM texts). We also explore the possibility of creating a proxy CFM: an LLM fine-tuned on a similar dataset to one used by the real ‘content farm’. We find that even a small amount of fine-tuning data suffices for creating a successful detector, but we need to know which base LLM is used, which is a major challenge. Our results suggest that there are

currently no practical methods for detecting synthetic news-like texts ‘in the wild’, while generating them is too easy. We highlight the urgency of more NLP research on this problem.

This paper has been awarded a Senior Area Chair award at ACL 2024.

DOI: 10.18653/v1/2024.acl-long.817

Using Large Language Models to Create Narrative Events

V. Bartalesi, E. Lenzi, C. De Martino

PeerJ Computer Science, vol. 10. PeerJ, 2024.

Narratives play a crucial role in human communication, serving as a means to convey experiences, perspectives, and meanings across various domains. They are particularly significant in scientific communities, where narratives are often utilized to explain complex phenomena and share knowledge. This article explores the possibility of integrating large language models (LLMs) into a workflow that, exploiting the Semantic Web technologies, transforms raw textual data gathered by scientific communities into narratives. In particular, we fo-

cus on using LLMs to automatically create narrative events, maintaining the reliability of the generated texts. The study provides a conceptual definition of narrative events and evaluates the performance of different smaller LLMs compared to the requirements we identified. A key aspect of the experiment is the emphasis on maintaining the integrity of the original narratives in the LLM outputs, as experts often review texts produced by scientific communities to ensure their accuracy and reliability. We first perform an evaluation on a corpus of five

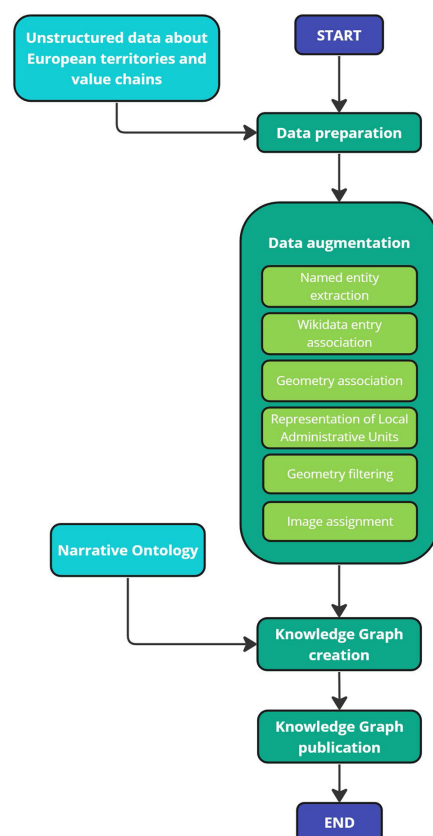
narratives and then on a larger dataset comprising 124 narratives. LLaMA 2 is identified as the most suitable model for generating narrative events that closely align with the input texts, demonstrating its ability to generate high-quality narrative events. Prompt engineering techniques are then employed to enhance the performance of the selected model, leading to further improvements in the quality of the generated texts.

DOI: 10.7717/peerj-cs.2242

A Semantic Knowledge Graph of European Mountain Value Chains

V. Bartalesi, G. Coro, E. Lenzi, N. Pratelli, P. Pagano, M. Moretti, G. Brunori

Scientific Data, vol. 11. Nature Publishing Group, 2024.



The United Nations forecast a significant shift in global population distribution by 2050, with rural populations projected to decline. This decline will particularly challenge mountain areas' cultural heritage, well-being, and economic sustainability. Understanding the economic, environmental, and societal effects of rural population decline is particularly important in Europe, where mountainous regions are vital for supplying goods. The present paper describes a geospatially explicit semantic knowledge graph containing information on 454 European mountain value chains. It is the first large-size, structured collection of information on mountain value chains. Our graph, structured through ontology-based

semantic modelling, offers representations of the value chains in the form of narratives. The graph was constructed semi-automatically from unstructured data provided by mountain-area expert scholars. It is accessible through a public repository and explorable through interactive Story Maps and a semantic Web service. Through semantic queries, we demonstrate that the graph allows for exploring territorial complexities and discovering new knowledge on mountain areas' environmental, societal, territory, and economic aspects that could help stem depopulation.

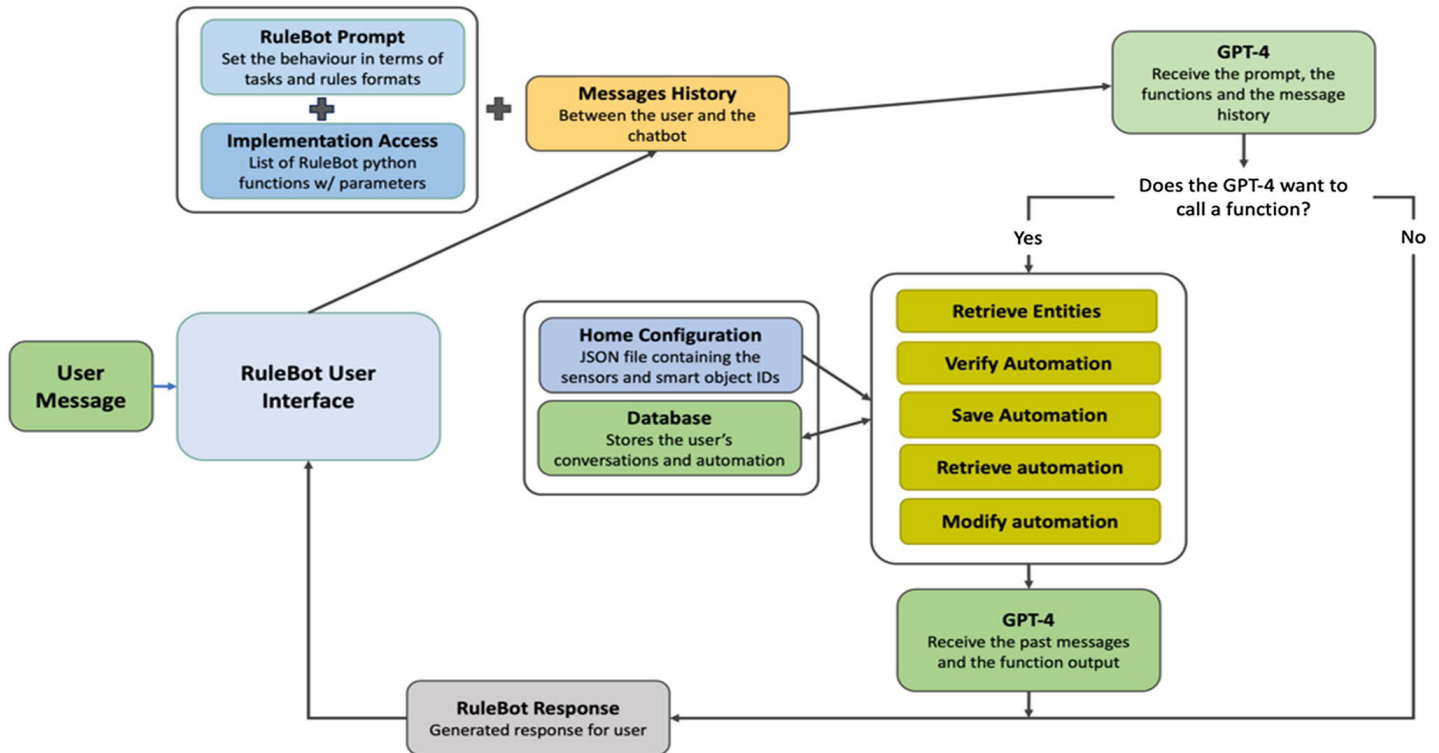
DOI: 10.1038/s41597-024-03760-9

Conceptual flowchart of our data preparation, augmentation, validation, and publication workflow.

A Conversational Agent for Creating Automations Exploiting Large Language Models

S. Gallo, F. Paternò, A. Malizia

Personal and Ubiquitous Computing, vol. 28. Springer, 2024.



The RuleBot++ architecture.

The proliferation of sensors and smart Internet of Things (IoT) devices in our everyday environments is reshaping our interactions with everyday objects. This change underlines the need to empower non-expert users to easily configure the behaviour of these devices to align with their preferences and habits. At the same time, recent advances in generative transformers, such as ChatGPT,

have opened up new possibilities in a variety of natural language processing tasks, enhancing reasoning capabilities and conversational interactions. This paper presents RuleBot++, a conversational agent that exploits GPT-4 to assist the user in the creation and modification of trigger-action automations through natural language. After an introduction to motivations and related

work, we present the design and implementation of RuleBot++ and report the results of the user test in which users interacted with our solution and Home Assistant, one of the most used open-source tools for managing smart environments.

DOI: 10.1007/s00779-024-01825-5

Extracting Mediterranean Hidden Fishing Hotspots Through Big Data Mining

G. Coro, L. Pavirani, A. Ellenbroek
IEEE Access, vol. 12. IEEE, 2024.

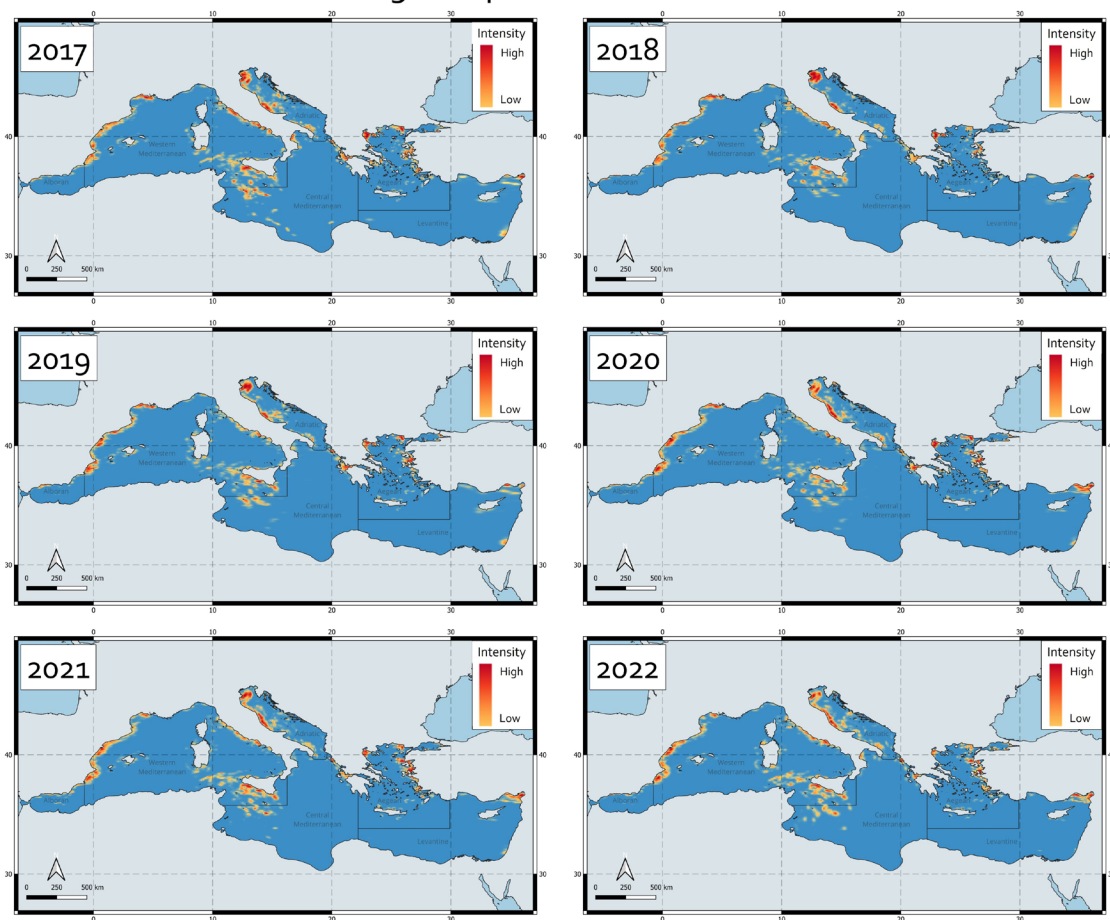
Monitoring fishing activity is crucial for fisheries management and governments to ensure regulatory compliance and sustainable marine ecosystems. Analysing vessel movements provides insights into fishing dynamics, aiding decision-making. Additionally, measuring unmonitored fishing activity (hidden fishing) helps counteract the underestimation of fishing pressure. Big data analysis can reveal fishing patterns and hidden activities from vessel position and speed data, such as those transmitted

by fleets carrying Automatic Identification Systems (AIS). We used an Open Science-compliant (reproducible, repeatable, and reusable) cloud computing-based big data analysis to estimate the manifest, total, and hidden fishing distributions of AIS-carrying vessels in the Mediterranean Sea from 2017 to 2022, processing about 1.6 billion vessel speed and position data. We estimated the principal hotspots of hidden fishing over the years and the potentially involved stocks from these data. We also assessed whether

the hotspots corresponded to illegal fishing or AIS communication issues and concluded that most hotspots potentially corresponded to illegal fishing. Our manifest fishing distribution agreed with another produced through machine learning by the Global Fishing Watch. We developed a fast and reusable approach that can produce new information to help management authorities understand the extent of hidden fishing.

DOI: 10.1109/ACCESS.2024.3416389

Hidden Fishing Hotspots of the Mediterranean Sea



Annual hidden fishing hotspots in the Mediterranean at 0.1° spatial resolution. A linear classification of the kernel density estimation range in each year was used for intensity representation.

Climate Change Effects on Animal Presence in the Massaciuccoli Lake Basin

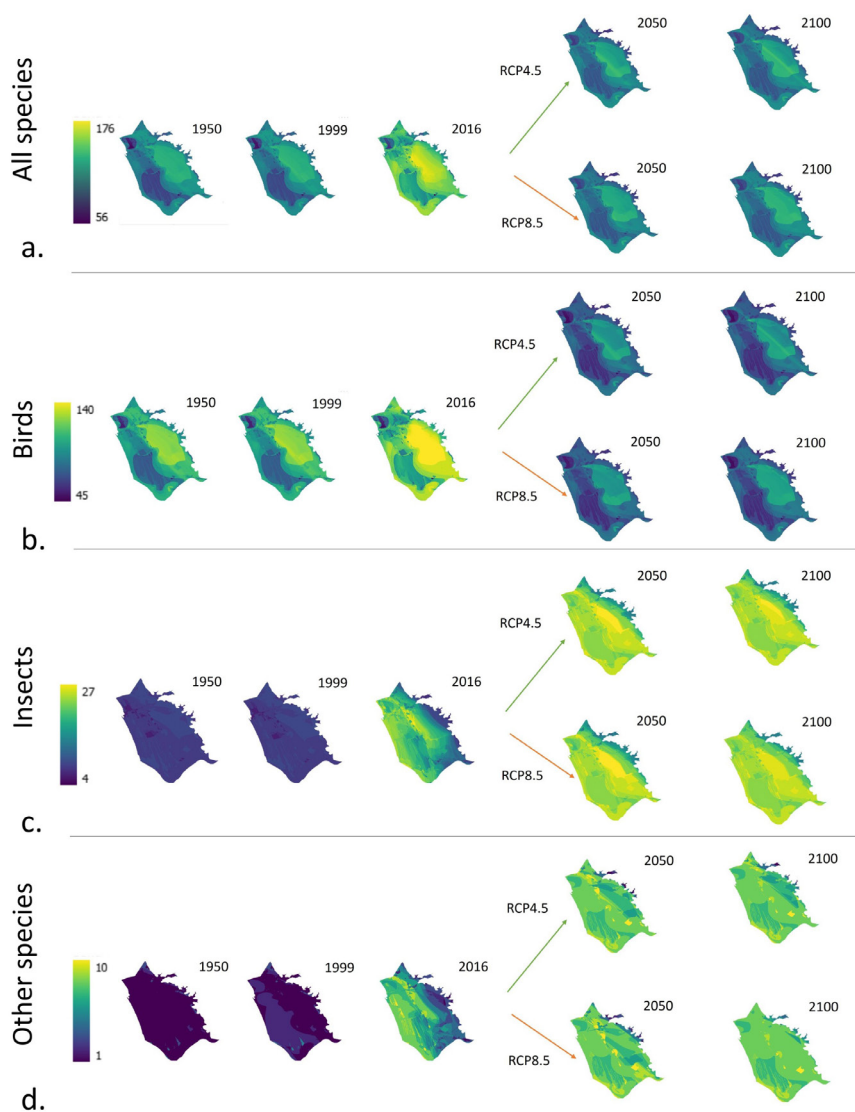
G. Coro, P. Bove, I. Baneschi, A. Bertini, L. Calvisi, A. Provenzale
Ecological Informatics, vol. 81. Springer, 2024.

Big-data mining approaches based on Artificial Intelligence models can help forecast biodiversity changes before they happen. These approaches can predict macroscopic species distribution patterns and trends that can inform preventive measures to avoid

the loss of ecosystem functions and services. They can, therefore, help study and mitigate climate change implications on biodiversity conservation in fragile ecosystems. Wetlands are particularly fragile ecosystems where climate change poses severe

risks and has dramatically reduced their size over the past century, with profound consequences on biodiversity and ecosystem services. Through big-data mining approaches, we can predict future wetland biodiversity trends in the context of climate change. This paper proposes such predictive analysis for a specific wetland: The Massaciuccoli Lake basin in Tuscany, Italy. This basin is a critical tourist attraction due to its rich biodiversity, making it an area of interest for citizens, tourists, and scientists. However, the region's suitability for native and non-native species is at risk due to climate and land-use change. Using machine-learning models, we predict the potential effects of climate change on animal spatial distribution in the basin under different greenhouse gas emission scenarios. The results suggest that habitat suitability has generally improved from 1950 to today, presumably owing to the targeted conservation strategies adopted in the area, but climate change will severely reduce bird biodiversity by 2050 while favouring several insect species' proliferation and other species' habitat change, even under a medium-emission scenario. This will lead to significant changes in the basin's biodiversity. Our methodology is adaptable to other wetland basins, being fully based on open data and models. The spatially explicit modelling used in this research provides valuable information for policymakers and spatial planners, complementing traditional biodiversity trend analyses.


DOI: 10.1016/j.ecoinf.2024.102644X



Species richness distributions over the years and the RCP4.5 and RCP8.5 scenarios for 4 different animal groups living in the Massaciuccoli Lake basin: All selected animal species (a), birds (b), insects (c), and other species (reptiles, amphibians, mammals, fishes and crustaceans) (d).

Fast, Interpretable, and Deterministic Time Series Classification with a Bag-of-Receptive-Fields

F. Spinnato, R. Guidotti, A. Monreale, M. Nanni
 IEEE Access, vol. 12. IEEE, 2024.



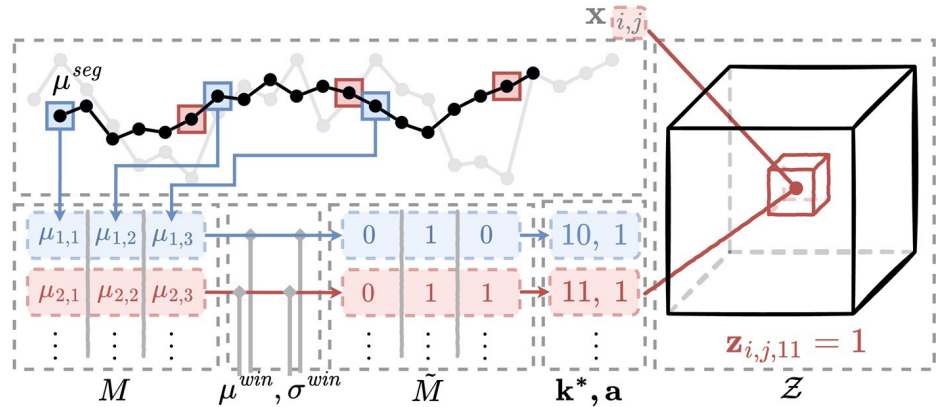
Two receptive fields (blue and red) with a length of $w=9$, divided into segments of length $q=3$. Dilation is the number of steps between consecutive observations within the receptive field. $d \cdot q$ is the segment hop, i.e., the number of steps between consecutive segments. The parameter $s=5$ is the stride, i.e., the distance between consecutive receptive fields.

The current trend in the literature on Time Series Classification is to develop increasingly accurate algorithms by combining multiple models in ensemble hybrids, representing time series in complex and expressive feature spaces, and extracting features from different representations of the same time series. As a consequence of this focus on predictive performance, the best time series classifiers are black-box models, which are not understandable from a human standpoint. Even the approaches that are regarded as interpretable, such as shapelet-based ones, rely on randomization to maintain computational efficiency. This poses challenges for interpretability, as the explanation can change from run to run. Given these limitations, we propose the Bag-Of-Receptive-Field (BORF), a fast, interpretable, and deterministic time series transform. Building upon the classical Bag-Of-Patterns, we bridge the gap between convolutional operators and discretization,

enhancing the Symbolic Aggregate Approximation (SAX) with dilation and stride, which can more effectively capture temporal patterns at multiple scales. We propose an al-

gorithmic speedup that reduces the time complexity associated with SAX-based classifiers, allowing the extension of the Bag-Of-Patterns to the more flexible Bag-Of-Receptive-Fields, represented as a sparse multivariate tensor. The empirical results from testing our proposal on more than 150 univariate and multivariate classification datasets demonstrate good accuracy and great computational efficiency compared to traditional SAX-based methods and state-of-the-art time series classifiers, while providing easy-to-understand explanations.

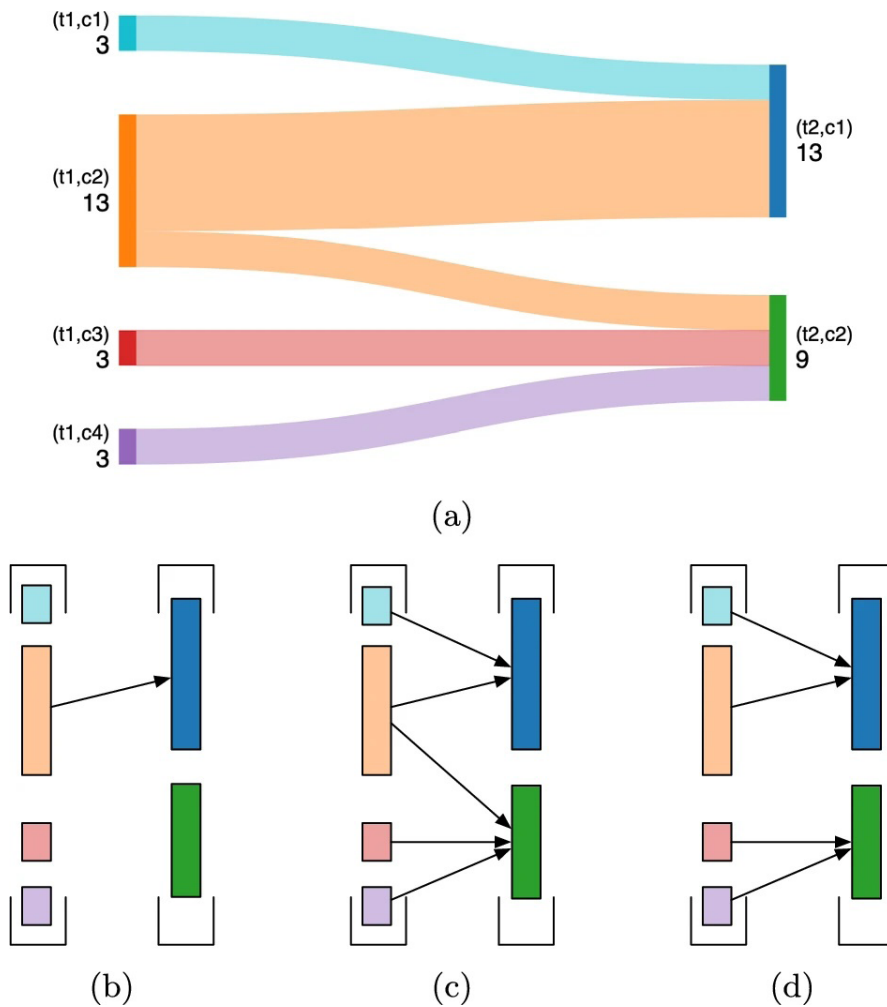
DOI: 10.1109/ACCESS.2024.3464743



A simplified schema of Algorithm 2 for the two receptive fields of Figure 1, extracted from signal $X_{i,j}$. First, the moving average μ^{seg} is computed. Values in μ^{seg} are used to fill M as in Equation (9). The two receptive fields are then normalized, binned into the SAX words $[0,1,0]$ and $[0,1,1]$, and hashed into the integers and respective counts $10,1$ and $11,1$. This allows the update of \mathcal{Z} . E.g., for the red word, $z_{i,j,11}=1$.

Describing Group Evolution in Temporal Data Using Multi-Faceted Events

A. Failla, R. Cazabet, G. Rossetti, S. Citraro
Machine Learning, vol. 113. Springer, 2024.



A realistic group evolution scenario (a), with corresponding event graph using different matching functions: b Jaccard coefficient, high threshold (Kalnis et al., 2005), c Jaccard coefficient, low threshold (Kalnis et al., 2005), d intersection over minimum size, as defined in Hopcroft et al. (2004)

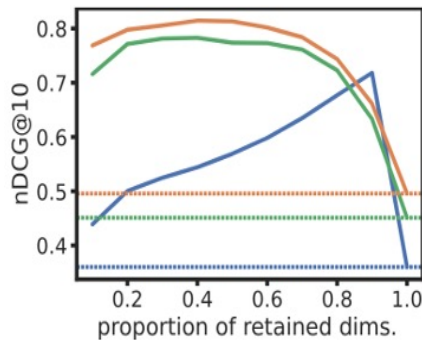
Groups—such as clusters of points or communities of nodes—are fundamental when addressing various data mining tasks. In temporal data, the predominant approach for characterizing group evolution has been through the identification of “events”. However, the events usually described in the literature, e.g., shrinks/growths, splits/merges, are often arbitrarily defined, creating a gap between such theoretical/predefined types and real-data group observations. Moving beyond existing taxonomies, we think of events as “archetypes” characterized by a unique combination of quantitative dimensions that we call “facets”. Group dynamics are defined by their position within the facet space, where archetypal events occupy extremities. Thus, rather than enforcing strict event types, our approach can allow for hybrid descriptions of dynamics involving group proximity to multiple archetypes. We apply our framework to evolving groups from several face-to-face interaction datasets, showing it enables richer, more reliable characterization of group dynamics with respect to state-of-the-art methods, especially when the groups are subject to complex relationships. Our approach also offers intuitive solutions to common tasks related to dynamic group analysis, such as choosing an appropriate aggregation scale, quantifying partition stability, and evaluating event quality.

DOI: 10.1007/s10994-024-06600-4

Dimension Importance Estimation for Dense Information Retrieval

G. Faggioli, N. Ferro, R. Perego, N. Tonellotto

Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24). ACM, 2024.



nDCG@10 on ROBUST04 as a function of the fraction of dimensions used for computing query-document similarity.

Recent advances in Information Retrieval have shown the effectiveness of embedding queries and documents in a latent high-dimensional space to compute their similarity. While operating on such high-di-

mensional spaces is effective, in this paper, we hypothesize that we can improve the retrieval performance by adequately moving to a query-dependent subspace. More in detail, we formulate the Manifold Clustering (MC) Hypothesis: projecting queries and documents onto a subspace of the original representation space can improve retrieval effectiveness. To empirically validate our hypothesis, we define a novel class of Dimension IMportance Estimators (DIME). Such models aim to determine how much each dimension of a high-dimensional representation contributes to the quality of the final ranking and provide an empirical method to select a subset of dimensions where to project the query and the documents. To support our hypothesis, we propose an oracle DIME, capable of effectively selecting dimensions and almost doubling the retrieval

performance. To show the practical applicability of our approach, we then propose a set of DIMEs that do not require any oracular piece of information to estimate the importance of dimensions. These estimators allow us to carry out a dimensionality selection that enables performance improvements of up to +11.5% (moving from 0.675 to 0.752 nDCG@10) compared to the baseline methods using all dimensions. Finally, we show that, with simple and realistic active feedback, such as the user's interaction with a single relevant document, we can design a highly effective DIME, allowing us to outperform the baseline by up to +0.224 nDCG@10 points (+58.6%, moving from 0.384 to 0.608).

DOI 10.1145/3626772.3657691

Efficient Inverted Indexes for Approximate Retrieval over Learned Sparse Representations

S. Bruch, F.M. Nardini, C. Rulli, R. Venturini

Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24). ACM, 2024.

Learned sparse representations form an attractive class of contextual embeddings for text retrieval. That is so because they are effective models of relevance and are interpretable by design. Despite their apparent compatibility with inverted indexes, however, retrieval over sparse embeddings remains challenging. That is due to the distributional differences between learned embeddings and term frequency-based lexical models of relevance such as BM25. Recognizing this challenge, a great deal of research has gone into, among other things, designing retrieval algorithms tailored to the properties

of learned sparse representations, including approximate retrieval systems. In fact, this task featured prominently in the latest BigANN Challenge at NeurIPS 2023, where approximate algorithms were evaluated on a large benchmark dataset by throughput and recall. In this work, we propose a novel organization of the inverted index that enables fast yet effective approximate retrieval over learned sparse embeddings. Our approach organizes inverted lists into geometrically-cohesive blocks, each equipped with a summary vector. During query processing, we quickly determine if a block must be


evaluated using the summaries. As we show experimentally, single-threaded query processing using our method, Seismic, reaches sub-millisecond per-query latency on various sparse embeddings of the MS MARCO dataset while maintaining high recall. Our results indicate that Seismic is one to two orders of magnitude faster than state-of-the-art inverted index-based solutions and further outperforms the winning (graph-based) submissions to the BigANN Challenge by a significant margin.

DOI 10.1145/3626772.3657769

Harnessing Topological Machine Learning in Raman Spectroscopy: Perspectives for Alzheimer's Disease Detection via Cerebrospinal Fluid Analysis

F. Conti, M. Banchelli, V. Bessi, C. Cecchi, F. Chiti, S. Colantonio, C. D'Andrea, M. de Angelis, D. Moroni, B. Nacmias, M.A. Pascali, S. Sorbi, P. Matteini

Journal of the Franklin Institute, vol. 361. Elsevier, 2024.



Scheme of the topological machine learning pipeline. Starting from the data we produce a persistence diagram by means of a filtration, which is then vectorized through various methods and such vectors enter a machine learning algorithm which returns a classification with a certain accuracy.

The cerebrospinal fluid of 21 subjects who received a clinical diagnosis of Alzheimer's disease (AD) as well as of 22 pathological controls has been collected and analysed by Raman spectroscopy (RS). We investigated

whether the Raman spectra could be used to distinguish AD from controls, after a pre-processing procedure. We applied machine learning to a set of topological descriptors extracted from the spectra, achieving a high

classification accuracy of 86%. Our experimentation indicates that RS and topological analysis may be a reliable and effective combination to confirm or disprove a clinical diagnosis of Alzheimer's disease. The following steps will aim at leveraging the intrinsic interpretability of the topological data analysis to characterize the AD subtypes, e.g. by identifying the bands of the Raman spectrum relevant for AD detection, possibly increasing and/or confirming the knowledge about the precise molecular events and biological pathways behind the Alzheimer's disease.

DOI: 10.1016/j.jfranklin.2024.107249

Environmental Study and Stress-Related Biomarkers Modifications in a Crew During Analog Astronaut Mission EMMPOL 6

T. A. Giacom, S. Mrakic-Spota, G. Bosco, A. Vezzoli, C. Dellanoce, M. Campisi, M. Narici, M. Paganini, B. Foing, A. Kołodziejczyk, M. Martinelli, S. Pavanello

European Journal of Applied Physiology. Springer, 2024.

Human presence in space is increasingly frequent, but we must not forget that it is a hostile environment. We aimed to study the characteristics of experimental scenarios, to obtain data on human response to isolation, disruption of circadian rhythm and high levels of psychophysical stress.

In these experiments, we evaluated stress response in five young healthy subjects inside an earth-based moon-settlement-like habitat during a 1-week long analog astronaut mission. Wearable devices were used

to monitor daily step count of the subjects, physical activity, heart rate during physical exercise and at rest, and sleep parameters. From saliva and urine samples collected every day at awakening, we studied oxy-inflammation biomarkers and hormones (stress and appetite) were studied too.

At the end of the week, all subjects revealed an increase in oxidative stress and cortisol levels but no inflammation biomarkers variations, in conjunction with increasing time/daily exercise. Furthermore, a significant de-

crease in hours of sleep/day, sleep quality, and REM phase of sleep was recorded and correlated with the increase of reactive oxygen species.


Oxidative stress increased in a short period of time and may be attributed to the influence of psychological stress during confinement, as well as increased exercise and decreased amount of sleep. On a long-term basis, this could impact performance.

DOI: 10.1007/s00421-024-05575-3

ANN Uncertainty Estimates in Assessing Fatty Liver Content from Ultrasound Data

G. Del Corso, M.A. Pascali, C. Caudai, L. De Rosa, A. Salvati, M. Mancini, L. Ghiadoni, F. Bonino, M.R. Brunetto, S. Colantonio, F. Faita


Computational and Structural Biotechnology Journal, vol. 24. Elsevier, 2024.



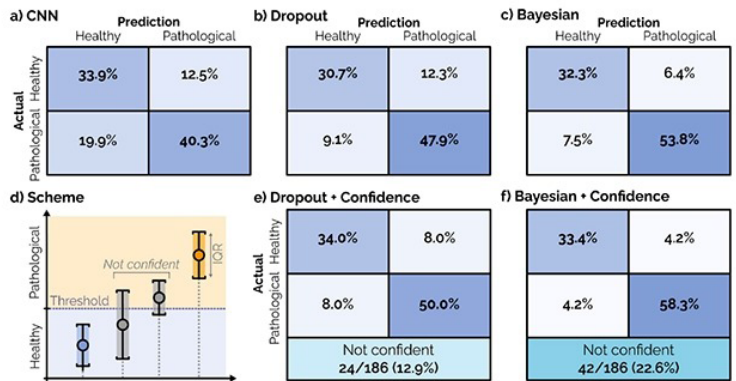

Comparison of three different convolutional architectures for assessing fatty liver content from US data:

1. The standard model (**deterministic CNN**)
2. The **MC Dropout**
3. The **Bayesian model**

with the goal of not only comparing the goodness of the predictions of the models, but also to have access to an evaluation of the level of uncertainty with which these architectures produce their output.



Graphical abstract.



This article uses three different probabilistic convolutional architectures applied to ultrasound image analysis for grading Fatty Liver Content (FLC) in Metabolic Dysfunction Associated Steatotic Liver Disease (MASLD) patients. Steatosis is a new silent epidemic and its accurate measurement is an impelling clinical need, not only for hepatologists, but also for experts in metabolic and cardiovascular diseases. This paper aims to provide a robust comparison between different uncertainty quantification strategies to identify advantages and drawbacks in a real clinical setting. We used a classical Convolutional Neural

Network, a Monte Carlo Dropout, and a Bayesian Convolutional Neural Network with the goal of not only comparing the goodness of the predictions, but also to have access to an evaluation of the uncertainty associated with the outputs.

We found that even if the prediction based on a single ultrasound view is reliable (relative RMSE [5.93%-12.04%]), networks based on two ultrasound views outperform them (relative RMSE [5.35%-5.87%]). In addition, the results show that the introduction of a “not confident” category contributes to increase the percentage of correctly

predicted cases and to decrease the percentage of mispredicted cases, especially for semi-invasive methods.


The possibility of having access to information about the confidence with which the network produces its outputs is a great advantage, both from the point of view of physicians who want to use neural networks as computer-aided diagnosis, and for developers who want to limit overfitting and obtain information about dataset problems in terms of out-of-distribution detection.

DOI: 10.1016/j.csbj.2024.09.021

Alignable Lamella Gridshells

D. Pellis

ACM Transactions on Graphics, vol. 43. ACM, 2024.




This paper presents a computational pipeline for the inverse design of alignable lamella gridshells, which are deployable grid structures capable of collapsing into a planar curved strip. These structures can find applications in architecture, facilitating transportation and installation processes.

Alignable lamella gridshells are 3D grid structures capable of collapsing into a planar strip. This feature significantly simplifies on-site assembly and ensures compactness for efficient transport and storage. However, designing these structures still remains a challenge. This paper tackles the inverse design problem of alignable lamella gridshells leveraging concepts from differential geom-

etry and Cartan's theory of moving frames. The study unveils that geodesic alignable gridshells, where lamellae are disposed tangentially to the surface, are limited to forming shapes isometric to surfaces of revolution. Furthermore, it demonstrates that alignable gridshells with lamellae arranged orthogonally to a surface can be realized only on a specific class of surfaces that meet

a particular curvature condition along their principal curvature lines. Finally, drawing on these theoretical findings, this work introduces novel computational tools tailored for the design of these structures.

DOI: 10.1145/3687898



Mannheim Multihalle': The structure shapes a Chebyshev net. The deployment alters the curvature of the beams and the angles of the net. The final shape is achieved by securing the boundary in a predetermined position (CC BY 3.0 H. Berberich). (b) A lamella gridshell based on an asymptotic Chebyshev net, from [Liu et al. 2023]. Such a net exists only on surfaces with constant negative Gaussian curvature (K-surfaces). (c) A geodesic alignable net with rotational symmetry, from [Tellier 2022a]. (d) 'Kinetic Umbrella' by Jonas Schikore and Eike Schling. This mechanism leverages an asymptotic alignable net with rotational symmetry (© J. Schikore).

Capacitive Touch Sensing on General 3D Surfaces

G. Palma, N. Pourjafarian, J. Steimle, P. Cignoni
ACM Transactions on Graphics, vol. 43. ACM, 2024.

Mutual-capacitive sensing is the most common technology for detecting multi-touch, especially on flat and simple curvature surfaces. Its extension to a more complex shape is still challenging, as a uniform distribution of sensing electrodes is required for consistent touch sensitivity across the surface. To overcome this problem, we propose a method to adapt the sensor layout of common capacitive multi-touch sensors to more complex 3D surfaces, ensuring high-resolution, robust multi-touch detection. The method automatically computes a grid of transmitter and receiver electrodes with as regular distribution as possible over a general 3D shape. It starts with the computation of a proxy geometry by quad meshing used to place the electrodes through the dual-edge graph. It then arranges electrodes on the surface to minimize the number of touch controllers required for capacitive sensing and the number of input/output pins to connect the electrodes with the controllers. We reach these objectives using a new simplification and clustering algorithm for a regular quad-patch layout. The reduced patch layout is used to optimize the routing of all the structures (surface grooves and internal pipes) needed to host all electrodes on the surface and inside the object's volume, considering the geometric constraints



Our computation fabrication method for multi-touch sensing on general 3D surfaces. (Left) 3D models with grooves to accommodate the touch sensor conductors and the internal pipes to connect the surface lines with the touch controller. (Center) Photo of the 3D printed prototypes equipped with the touch sensor grid made of enamelled unipolar solid copper conductor. (Right) An example of touch interaction with two fingers with the relative positions computed on the 3D model.

of the 3D shape. Finally, we print the 3D object prototype ready to be equipped with the electrodes. We analyze the performance of the proposed quad layout simplification and clustering algorithm using different quad meshing and characterize the signal quality and accuracy of the capacitive touch sensor for different non-planar geometries.

The tested prototypes show precise and robust multi-touch detection with good Signal-to-Noise Ratio and spatial accuracy of about 1mm.

DOI: 10.1145/3658185.



Processing pipeline. Starting from a triangle mesh, we compute a proxy quad mesh, cluster the quads in a coarse quad patch layout, and pack this layout in the sensing regions of the touch controllers. Then, we compute the final geometry of the prototype, generating the surface grooves to accommodate the sensor conductors and internal pipes to connect these conductors with the controllers. The final step is the physical fabrication of the prototype.

Seismic Response and Ambient Vibrations of a Medieval Tower in the Mugello Area (Italy)

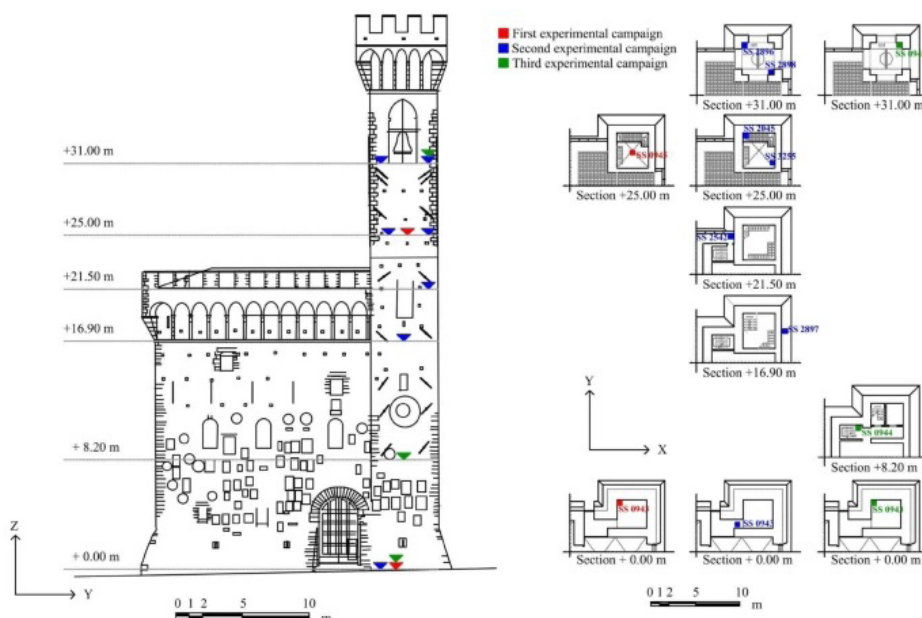
R.M. Azzara, V. Cardinali, M. Girardi, C. Padovani, D. Pellegrini, M. Tanganelli
Journal of Civil Structural Health Monitoring. Springer, 2024.

This paper describes the experimental campaigns on the Tower of the Palazzo dei Vicari in Scarperia, a village in the Mugello area (Tuscany) exposed to high seismic hazards. The first campaign was carried out from December 2019 to January 2020, and the Tower underwent the so-called Mugello seismic sequence, which featured an M 4.5 earthquake. Other ambient vibration tests were repeated in June 2021 and September 2023 when another seismic sequence struck the area near Scarperia. These tests aimed to characterise the Tower's dynamic behaviour under ambient and seismic excitations and check the response of the Tower over time. The experimental results were then used to calibrate a finite-element model of the Tower and estimate its seismic vulnerability. Several numerical simulations were conducted on the calibrated model using the NOSA-ITACA code for nonlinear structural analysis of masonry buildings. The dynamic behaviour of the Tower subjected to a seismic sequence recorded in 2023 by a seismic station at the base was investigated by comparing the velocities recorded along the Tower's height with their numerical counterparts. Furthermore, several pushover analyses were conducted to investigate the collapse of the Tower as the load's distribution and direction varied.

DOI: 10.1007/s13349-024-00824-7



The historic centre of Scarperia, Tuscany: aerial view.



Geometry of the Tower (length in metres) and experimental setup of the monitoring campaign. Red: position of the seismometers during the first campaign (2019); blue: position during the second campaign (2021); green: position during the third campaign (2023).

A Configurable Software Model of a Self-Adaptive Robotic System

A software model for ProFeat, a model-checking analysis tool

Self-adaptive systems (SASs) often operate in dangerous and dynamic environments where human supervision is limited or impossible, such as under water or in space. It is therefore important to ensure that safety properties are maintained throughout system operations. Once in operation, SASs are frequently reconfigured, which means switching between different system configurations during runtime. The separate analysis of all these different yet partially redundant configurations is a tedious, time-consuming, and error-prone task, especially because it does not take into account the changes between configurations.

In “Formal Modelling and Analysis of a Self-Adaptive Robotic System” (published in LNCS 14300, Springer, 2023, https://doi.org/10.1007/978-3-031-47705-8_18), Päßler et al. highlighted the advantages of modelling such an SAS as a family of systems, where each family member corresponds to a possible configuration, which allows for family-based modelling and analysis

as a means to combat redundancy. Formal models and tools from the field of software product lines (SPLs) were used, combined with the knowledge that SASs can be implemented using a two-layered approach. This breaks the system down into a managed and a managing subsystem, where the managed subsystem implements the domain concerns (e.g., navigating a robot to a specific position), while the managing subsystem implements the adaptation logic (e.g., reconfiguring due to changing environmental conditions). This separation of concerns is catered for by ProFeat (<https://pchrson.github.io/profeat/>), a tool for probabilistic family-based model checking. ProFeat provides a means to simultaneously analyse, in one single run, a family of models, each corresponding to a valid configuration.

This resulted in a configurable software model of a self-adaptive robotic system, namely an Autonomous Underwater Vehicle (AUV) used to search for and follow a pipeline located on a seabed, which was written

by Juliane Päßler, Maurice H. ter Beek, Ferruccio Damiani, Einar Broch Johnsen, and S. Lizeth Tapia Tarifa and published as an original software publication in Science of Computer Programming, volume 240 (2024), <https://doi.org/10.1016/j.scico.2024.103221>. This paper illustrates how to perform analyses of such models using ProFeat, and how to modify and extend the models. The published software model of the AUV case study is relevant for both new and existing questions in research on SASs and dynamic SPLs, as well as in industry. In fact, during his SPLC 2023 keynote address, Kentaro Yoshimura, Chief Researcher at Hitachi, presented the use of dynamic SPLs for autonomous robotic systems as a new industrial challenge.

Contact: Maurice ter Beek, FMT Lab
maurice.terbeek@isti.cnr.it
<http://fmt.isti.cnr.it/~mtbeek/>

Code metadata description	
Current code version	v1.1.2
Permanent link to code/repository used for this code version	https://github.com/remaro-network/auv_profeat/releases/tag/SCP-2024
Permanent link to reproduce capsule	https://doi.org/10.5281/zenodo.13946884
Legal code license	Apache License 2.0
Coder versioning system used	git
Software code languages, tools and services used	PRISM input and property language, ProFeat, iFM23 virtual machine
Compilation requirements, operating environments and dependencies	Linux, Windows, MacOS (Intel for capsule)
If available, link to developer documentation/manual	https://github.com/remaro-network/auv_profeat/blob/scp-ifm_artifact/README.md
Support email for questions	julipas@uio.no

Metadata of the configurable software model

Best Paper Runner-up Award at SIGIR 2024

At the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, Washington D.C., USA, July 14-17, 2024



The paper “Efficient Inverted Indexes for Approximate Retrieval over Learned Sparse Representations” by Sebastian Bruch, Franco Maria Nardini, Cosimo Rulli, and Rossano Venturini won the Best Paper Runner-up Award.

Learned sparse representations are an attractive class of contextual embeddings for text retrieval because they combine effectiveness in relevance modeling with interpretability by design. Despite their compatibility with inverted indexes, retrieval over sparse embeddings poses challenges due to distributional differences between these embeddings and traditional term frequency-based models like BM25. This challenge has involved designing retrieval algorithms tailored to the properties of learned sparse representations, including approximate retrieval systems. This task featured prominently in the latest BigANN Challenge at NeurIPS 2023, where approximate algorithms were evaluated on a large benchmark dataset by throughput and recall. We propose a novel organization of the inverted index to enable fast and effective approximate retrieval over learned sparse embeddings. The approach organizes inverted lists into geometrically cohesive blocks, each summarized by a vec-

tor. During query processing, the summaries allow for rapid determination of whether a block needs evaluation. Experiments demonstrate that our novel approach, called “Seismic”, achieves sub-millisecond query latency on various sparse embeddings of the MS MARCO dataset while maintaining high recall. It is one to two orders of magnitude faster than state-of-the-art inverted index-based solutions and significantly outperforms graph-based methods that won the BigANN Challenge at NeurIPS 2023.

This groundbreaking work provides a significant contribution to the field of information retrieval, combining efficiency with high recall to meet the demands of modern data-intensive applications.

Contact: Franco Maria Nardini, HPC Lab
francomaria.nardini@isti.cnr.it

Best Paper Honorable Mention at ACM SIGGRAPH 2024

At the SIGGRAPH 2024 Conference, Denver, USA, July 28 – August 1, 2024



The paper “Capacitive Touch Sensing on General 3D Surfaces” by Gianpaolo Palma, Narges Pourjafarian, Jürgen Steimle, and Paolo Cignoni received an Honorable Mention. Selected from approximately 840 submissions, the paper introduces a pioneering computational method for adapting capacitive touch sensing to complex 3D shapes. This method achieves a touch position accuracy of approximately 1mm (<https://vcg.isti.cnr.it/publication/2024/PPSC24-3dTouch/>).

Contact: Gianpaolo Palma, VC Lab
gianpaolo.palma@isti.cnr.it
<https://doi.org/10.1145/3658185>

Best Talk Award at CCS2024

At the Conference on Complex Systems 2024 in Exeter, UK, September 2-6, 2024

The paper “Modeling Gentrification as a Relocation Flow-Based, Spatio-Temporal Process” by Giovanni Mauro, Nicola Pedreschi, Luca Pappalardo, and Renaud Lambiotte received the Best Talk Award.

It also proposes new measures for tracking gentrification, including one that acts as an early warning indicator. The findings show that more densely populated cities face higher risks of gentrification.

Selected from approximately 500 submissions, the paper demonstrates how different income groups drive gentrification.

Contact: Giovanni Mauro, KDD Lab
giovanni.mauro@isti.cnr.it
<https://ccs24.cssociety.org/>



Best Paper Award at ASONAM 2024

At the 16th International Conference on Advances in Social Networks Analysis and Mining, Rende (CS), Italy, September 2-6, 2024

The paper entitled “Beyond Boundaries: Capturing Social Segregation on Hypernetworks” by Francesco Cauteruccio (University of Salerno), Andrea Failla and Giulio Rossetti (Cnr-Isti) won the Best PaperAward.

Complex social systems have gained momentum due to renewed interest in higher-order topologies, which has led to the emergence of hypernetwork science.

A key phenomenon in these social systems is segregation—the extent to which entities within the network are separated or clustered based on specific attributes or features. Beyond Boundaries introduces a new framework for studying segregation within hypernetworks.

This work extends traditional segregation measures, initially developed for dyadic (pairwise) network structures, to polyadic (multi-entity) structures. Additionally, it



introduces a new measure, Random Walk HyperSegregation (RWHS), which employs random walkers to estimate segregation at multiple scales.

limitations of classical segregation measures when applied to higher-order topologies.

Contact: Giulio Rossetti, KDD Lab
giulio.rossetti@isti.cnr.it

Through extensive experiments on both synthetic and online social media data, Beyond Boundaries demonstrates the effectiveness of RWHS and highlights the

Cor Baayen Young Researcher Award 2024



Nicola Messina (Cnr-Isti, on the left) and Björn Levin (ERCIM President, on the right)

Nicola Messina, 31, researcher at CNR-ISTI, was awarded the Cor Baayen Young Research Award.

This award is organized annually by the European Research Consortium for Informatics and Mathematics (ERCIM). It acknowledges the research work of young talent in computer science and applied mathematics. The award was presented during the consortium's autumn meeting, where Nicola presented his research.

Nicola's research is distinguished by its high quality, interdisciplinary approach, and significant impact. His work spans artificial intelligence, computer vision, deep learning, and multimedia information retrieval. He has made substantial contributions to both the theoretical and applied aspects of these fields, with measurable scientific and practical outcomes.

During his PhD, Nicola began exploring the ability of neural networks to understand and process relationships between objects in computer vision. He boosted multimodal artificial intelligence by developing innovative and efficient methods for aligning representations of complex neural networks in both visual and language modalities.

His work in multimedia information retrieval has demonstrated immediate real-world applicability. This is exemplified by VISIONE, a large-scale video search system that won the 2024 international Video Browser Show-down competition and secured second place in 2023. VISIONE has been employed by the Italian national public broadcaster RAI, as part of the AI4Media European project, to facilitate efficient browsing of audiovisual archives. More recently, Nicola has been interested in spotting critical limitations in vision-language models for fine-grained understanding. He has expanded his research focus to include the application of atten-

tive deep learning techniques for structural health monitoring and the preservation of cultural heritage.

"Nicola Messina is one of the young very talented researchers of our institute, which this award is testament to. It's a prestigious recognition that makes us all particularly proud, especially in the light of the excellent performance by the National Research Council in last year's competition, where top European research institutes compete each year", says Roberto Scopigno, director of CNR-ISTI.

Nicola's professional experience also includes international collaborations with esteemed European universities and participation in various Italian and European research projects, such as AI4Media, AI4EU, INAROS, AI4CHSites, ADA, and Smart News. In addition to his research, he is involved in teaching and dissemination activities, offering instruction and guidance to students on topics related to deep learning, computer vision, and multimodal processing.

In 2021 and 2022, Nicola received the ISTI Young Researcher Award, which recognized him as one of the top young researchers (under 32 years old) at CNR-ISTI.

Contact: Nicola Messina, AIMH Lab
nicola.messina@isti.cnr.it

Best Paper Award: Individual Alternative Routing based on Road Popularity

Received the Best Paper Award at the NetMob 2024 conference.



Giuliano Cornacchia and Luca Pappalardo (Cnr-Isti) won the Best Paper Award at the NetMob2024 Conference.

The NetMob2024 Conference (<https://netmob.org/>) is the leading conference on the analysis of mobile phone datasets in social, urban, societal and industrial problems.

Overview of the paper

The paper introduces an innovative routing algorithm that distributes vehicles across road networks to minimize CO2 emissions. Unlike conventional methods, which often overlook road popularity—the frequency with which certain roads are used across different city zones—this new approach incorporates a metric to measure and account for this popularity. By dynamically adjusting routing recommendations, the algorithm distributes traffic more evenly throughout the network, preventing the overloading of popular roads. Through simulations in Florence, Milan, and Rome, our solution showed a notable reduction in CO2 emissions, achieving up

to a 23.57% decrease compared to other state-of-the-art algorithms. Additionally, it enhanced travel smoothness by reducing encounters with regulated intersections, thus promoting a more fluid driving experience.

Overall, this approach offers a sustainable alternative to urban mobility by reducing congestion and emissions, while providing drivers with efficient route options that avoid highly congested areas. This dynamic and adaptive routing method represents a significant step forward in environmentally conscious traffic management.

Contact: G. Cornacchia, L. Pappalardo,
KDD Lab
giuliano.cornacchia@isti.cnr.it
luca.pappalardo@isti.cnr.it
<https://netmob.org>

“Leonardo Lesmo” 2024 Award

Lorenzo Volpi has won the “Leonardo Lesmo” 2024 New Graduate Award from “Associazione Italiana per l'Intelligenza Artificiale” for his MSc thesis entitled “Predicting Classifier Accuracy under Prior Probability Shift.”

Volpi received his master's degree with honors in Computer Science from the University of Pisa, supervised by ISTI researchers Andrea Esuli, Alejandro Moreo,

and Fabrizio Sebastiani. Volpi addressed the problem of estimating the accuracy of classification algorithms on data with different distributions from the classifiers' training data. Solving this problem has important practical applications in the field of AI.

The proposed solution is based on using quantification algorithms and is at the vanguard of the state of the art on the

problem. Volpi currently continues his research with a research grant at ISTI. The award was presented during the association's 23rd annual conference in Bolzano, Nov. 25-28.

Contact: Lorenzo Volpi, AIMH Lab
lorenzo.volpi@isti.cnr.it
<https://aixia2024.events.unibz.it/>

Best Poster Award, PhD Day 2024 of DII

At the School of Engineering in Pisa



A total of 34 first-year PhD students had the opportunity to present their work, exchange ideas, and share projects.

The participants themselves selected the best poster, and this year's winner is Laura Pavirani, a marine biologist and now a PhD student in Information Engineering.

As in previous years, the department's PhD students showcased their work in a poster session. This year's event, held at the School of Engineering, aimed to encourage participation and foster exchanges among PhD students, undergraduate students, and fac-

ulty members. Strong turnout was achieved thanks to the efforts of organizers Nicola Vanello, Chiara Magliaro, and Fulvio Gini, with support from the University of Pisa.

"The poster session," commented the organizers, "is, for many first-year PhD students, a 'debut into society,' as it is their first opportunity to discuss preliminary research results with an audience in the field. For a growing department spread across multiple sites, it's also a chance to meet colleagues, share first impressions of the PhD program, and build a sense of community. We also noticed great interest and curiosity from Master's students. The PhD students themselves selected the best work, encouraging them to explore all their colleagues' projects and ask questions."

This year's winner, Laura Pavirani, presented her poster "Interplay of Climate, Fishing, and Biodiversity: Risk Assessment in the Mediterranean Sea," a study on big data analysis with climate and ecological relevance.

"I decided to pursue a PhD in Information Engineering," she explained, "despite it being outside my academic background, to acquire advanced data analysis methodologies. These would allow me to gain a broader perspective on how to study, understand, and protect marine environments. My work proposes a methodology for identifying ar-

reas where the overlap of environmental conditions, climate change factors, fishing activities (both legal and illegal), and biodiversity creates situations potentially at high risk of environmental degradation and over-exploitation of marine resources. The innovation lies in solving a risk assessment issue through a combination of clustering and statistical techniques. This methodology identified high-risk areas in the Mediterranean, particularly in the shallow waters of the Tyrrhenian Sea, northern Adriatic, and Aegean Sea.

Other sensitive areas include the coasts of Spain and France, the Strait of Sicily, and eastern Turkey, where illegal fishing emerged as a prevalent factor. These maps are valuable tools for monitoring authorities such as the FAO, the General Fisheries Commission for the Mediterranean, and Mediterranean governments, to whom they will be presented soon. The research is part of the ITINERIS project of the PNRR and is funded by the Institute of Marine Sciences of the National Research Council in Lerici (CNR-ISMAR), in collaboration with the Institute of Information Science and Technologies 'A. Faedo' in Pisa (CNR-ISTI)."

Contact: Laura Pavirani, Infrascience Lab
laura.pavirani@isti.cnr.it

Grants for Young Mobility

The ISTI Grants for Young Mobility (GYM) program enables young researchers (under 34) to carry out research in cooperation with foreign Universities and Research Institutions of clear international standing.

This is a remarkable and important opportunity to broaden their expertise, enhance international collaboration, and contribute to cutting-edge scientific advancements. It complements similar CNR programs.

The reference person is Franco Maria Nardini (francomaria.nardini@isti.cnr.it).
The winners in 2024 are:

First Call 2024



Katherine E. Abramski



Daniele Gambetta



Chiara Mannari



Giovanni Puccetti

Second Call 2024



Andrea Failla

Indoor Localization Technologies for Cultural Heritage: the MVL app



An innovative support system for museum visitors, capable of automatically recognizing artworks being viewed, provides informational sheets and photographic details. This app, designed and developed by the Wireless Network Lab of the Institute of Science and Information Technologies "A. Faedo" (CNR-ISTI) in Pisa, was presented on October 18, 2024, at the Circolo dei Forestieri in Bagni di Lucca. The app is now ready for use by visitors at two fascinating museums: the Museum of Sacred Art in Benabbio and the Museum of San Cassiano di Controne, both located in

Val di Lima (Lucca). It is available for download under the name "Musei Val di Lima (MVL)" and can be used on Android and iOS mobile devices.

"The development of the app, based on Bluetooth technology, was made possible by advancements in short-range wireless technologies. These enable the identification of the position of people and objects in indoor environments where satellite localization (GPS) may be ineffective. The application of indoor localization systems is particularly useful in settings such as museums and his-

torical sites, supporting visitors during their explorations," explains Michele Girolami, a researcher at the Wireless Network Lab of CNR-ISTI.

The presentation event also featured the director of the Institute, Roberto Scopigno, who emphasized the importance of this project in strengthening the "connection between artwork and local area." During the presentation, Prof. Sonia Maffei, head of the LIMES Laboratory at the University of Pisa, explained how one of the project's goals is to foster cultural and social transformation.

"The digital world opens up an endless range of possibilities for experiencing museum works. With this project, technology becomes art," said Paola Antonella Andreuccetti, director of the Museum of San Cassiano di Controne. Together with Agnese Benedetti, head of the Museum of Sacred Art in Benabbio, she moderated the event. This experimental initiative aligns with one of CNR's missions: to share knowledge and make technologies widely accessible for the benefit of the community.

Contact: Michele Girolami, WN Lab
michele.girolami@isti.cnr.it

AI/ML Assisted 5G-NTN Integration for Optimized Network Slicing over Multi-Technology Architecture

(Time of Change)

At the 2024 edition of the Italian Networking Operators Group (ITNOG), Pietro Cassarà presented the initial results of an experiment applying reinforcement learning to manage routing in an IP-Radio Access Network (IP-RAN). The aim is to optimize the integration of terrestrial network technologies with non-terrestrial ones, such as LEO constellations, within the 5-6G technology framework. The experiment, is part

of the ENI Proof of Concept #22 proposed by the European Telecommunications Standards Institute, for which Pietro Cassarà is responsible. Contributions were made by Alberto Gotta of ISTI-CNR, Riccardo Burrai, and Emiljan Kolaj of SIRIUS Technology, and supported by Huawei thanks to Aldo Artigiani and Luca Vit. The activity is also part of the long-standing scientific collaboration between Huawei and ISTI, for which ISTI re-

ceived a hardware infrastructure donation to conduct research. The research is also managed thanks to recent contributions from Franca Debole, Andrea Dell'Amico, and Tommaso Piccioli.

Contact: Cassarà Pietro WN Lab
pietro.cassara@isti.cnr.it
<https://www.linkedin.com/feed/update/urn:li:activity:7251557355518050304/>

Marco Malvaldi Joins ISTI to Unlock the Potentials of LLMs for Writing



Fabrizio Falchi (Cnr-Isti, on the left), Marco Malvaldi (writer, in the center) and Roberto Scopigno (Cnr-Isti Director, on the right)

The exceptional writing capabilities of the latest chatbots and virtual assistants based on Large Language Models (LLMs) pose a significant question for anyone attempting to write today—a scientist, a writer, or a lover. Am I really writing a text that an AI would not have been able to produce?

We have enlisted the help of Marco Malvaldi, a chemist and writer, whose unique style and wit have earned him a significant following both in Italy and abroad, we aim to investigate how LLMs and, more broadly, approaches in natural language processing

and computational linguistics can be used to evaluate different styles, assist in writing, and contribute to the highly debated theme of authorship attribution.

We believe that a continuous interdisciplinary dialogue between computer scientists and those who have made writing their profession is indispensable for the development of these research activities. Marco Malvaldi is an ideal partner in this research due to his highly successful novelist track record and his scientific background.

The first result of this collaboration was recently presented (“You write like a GPT”, by A. Esuli, F. Falchi, M. Malvaldi, and G. Puccetti) at the Tenth Italian Conference on Computational Linguistics (CliC-it 2024), where we investigated how Raymond Queneau’s “Exercises in Style” and its Italian version written by Umberto Eco are evaluated by automatic methods for detecting artificially generated text.

Contact: Fabrizio Falchi, AIMH Lab
fabrizio.falchi@isti.cnr.it

SoBigData Academy

The European Research Platform for social mining self-learning



Information has been transformed from isolated facts into a vast reservoir of hidden wealth, becoming essential to our daily lives. Big Data revolutionizes fields from targeted advertising to personalized health-care. However, to harness its full potential, users must acquire and refine new skills for exploring this vast data landscape.

As a research infrastructure, the SoBigData Academy has been created to enhance knowledge and opportunities. This is not just another online learning platform. Integrated into the SoBigData RI, the European Research Infrastructure for Big Data and Social Mining, the Academy offers a unique educational experience. It is entirely free and accessible to everyone, regardless of background or expertise.

The SoBigData Academy bridges the gap between users and leading European universities, including the University of Pisa, Scuola Superiore Sant'Anna, King's College London, and ETH Zürich. This network fos-

ters valuable professional connections, potentially opening new career paths.

The curriculum, curated by distinguished researchers and professionals from European institutions, provides diverse perspectives and cutting-edge knowledge, ensuring a well-rounded and enriching learning experience. The Academy is designed to cater to various audiences, including students, doctoral candidates, researchers, industry professionals, and those simply curious about this field.

MOOCs (Massive Open Online Courses), with seamless enrollment and adaptive learning speed and complexity. This choice exploits the key features of MOOCs, such as accessibility and flexibility, allowing students to access high-quality courses without geographical or time constraints. The affordability of MOOCs makes advanced education in social mining and AI accessible to a broader audience, democratizing knowledge in these crucial fields. Self-learning allows stu-

dents to progress independently, balancing their studies with other commitments. The Academy offers beginner and expert paths to tailor the experience to individual knowledge and skill levels.

In addition, the interactive and engaging content of SoBigData Academy stimulates active and participatory learning, which is essential for mastering complex data science concepts. The self-learning paradigm empowers the user to take control of the educational journey, choosing from a wide range of courses, including:

Legal and Ethical Aspects of Data Science: Exploration of data science's ethical and legal considerations and the current EU regulations.

Database: An introduction to database design, management, and SQL for efficient data storage and retrieval.

Basic Python: Foundational programming

covering essential syntax, functions, and data structures for beginners.

Information Retrieval: Techniques for effective retrieval of relevant information from vast datasets, including search algorithms.

Data Analysis: Core techniques for data manipulation, statistical analysis, and visualization to inform data-driven decisions.

Data Mining & Machine Learning: Key data mining and machine learning methods, including classification, clustering, and predictive modelling.

Data Analysis with Spark: Big data analysis with Apache Spark, focusing on distributed processing and large-scale data handling.

Complex Network Analysis: Study complex networks using graph theory to analyze social systems.

Text Analytics: Techniques for extracting, analyzing, and interpreting insights from textual data using up-to-date tools.

Data Theory and Society: Examines data's role in shaping society, emphasizing theoretical frameworks and societal impacts.

Artificial Intelligence: Introduction to AI concepts and algorithms, from foundational principles to practical applications in various fields.

Neural Networks & Deep Learning: Comprehensive coverage of neural networks,

covering architectures, training methods, and applications in AI.

Reinforcement Learning: Study of reinforcement learning algorithms for decision-making and autonomous systems.

Data Visualization and Storytelling: Effective data visualization and storytelling to convey insights clearly and impactfully.

Additional courses are planned for 2025, expanding the curriculum through the Academy's extensive academic network of SoBigData. This is just the beginning: the SoBigData Academy is to offer a diverse range of courses to deepen the user's knowledge and provide the skills needed to thrive in a data-driven society.

Once the user completes a course, they will be awarded a certificate recognized across Europe. The SoBigData RI supports and promotes this certification, which proves the acquisition of new skills and knowledge. Given that acquiring knowledge should be an intrinsically motivating and likeable process, the SoBigData Academy uses various engaging learning tools to keep the user motivated. Each course The Academy incorporates engaging tools like interactive lessons, quizzes, forums, and gamified elements such as crosswords and memory games. Learners also gain access to Jupyter Notebooks for hands-on coding practice, bridging theory and real-world application.

The SoBigData Academy opened on November 28, 2024, with an initial seven courses.

A second set of courses will launch in early 2025. This initiative is promoted through the Géant network, ensuring high-quality training materials.

We invite everyone to explore this innovative learning environment and participate in the educational opportunities offered by SoBigData Academy.

Contact: Roberto Trasarti, KDD Lab

roberto.trasarti@isti.cnr.it

Sara Lelli, KDD Lab,

Learning Designer of the Academy

sara.elli@isti.cnr.it

<http://www.sobigdata.eu/academy>



Ph.D. Dissertations

A Bridge between Persistent Homology and Group Equivariant Non Expansive Operators: Theory and applications

Author: Francesco Conti, University of Pisa, Italy

Supervisors: Davide Moroni, Cnr-Isti, Patrizio Frosini, University of Bologna, Maria Antonietta Pascali, Cnr-Isti

Topological Data Analysis (TDA) is proving to be an excellent tool for shape analysis of digital data. The recently found synergy with artificial intelligence gave rise to Topological Machine Learning (TML), which aims to combine the expressive power of computational topology with the accuracy of machine learning to provide a comprehensive and automatic framework for data classification. The aim of this thesis is twofold: to develop current applications of TML in practical scenarios, with emphasis on the most overlooked aspects of its pipeline, and to connect the theory of TDA with a broader class of maps, the Group Equivariant Non-Expansive Operators (GENEOs). In the first part of this dissertation, we develop a pipe-

line to study digital data by means of TML in order to validate the practical aspects of our theory. We apply this pipeline to benchmark and experimental datasets, achieving state-of-the-art accuracies in biomedical scenarios. Moreover, we perform an empirical but extensive study of the stability of features arising from the various homological dimensions with respect to noise and points distribution in the persistence diagram. Such a comparison is novel in the TML literature, and our findings show that results coming from the concatenation of each homological dimension available are the best approach in the vectorization step. We later expand on the main concept of TDA, proving that the functor that computes persistence dia-

grams can be seen as a particular instance of GENEOs. The GENEO framework allows us to inject arbitrary equivariances into a machine learning setting and represents a new possible approach to neural network architecture. Next, we fully present the theory of GENEOs and their properties, such as convexity and concavity, under suitable assumptions. This thesis expands the GENEO theory with two new tools to define such operators, namely using symmetric functions and a characterization theorem of linear GENEOs between arbitrary functional spaces. Finally, we develop a new neural network architecture with GENEOs instead of neurons and show its potential in a couple of applications.

Quantifying and Mitigating the Impact of Vehicular Routing on the Urban Environment

Author: Giuliano Cornacchia, University of Pisa, Italy

Supervisors: Luca Pappalardo, Cnr-Isti and Scuola Normale Superiore Pisa, Mirco Nanni, Cnr-Isti

Urbanization pressures cities to efficiently accommodate the increasing demand for mobility, making traffic optimization challenging due to the complex interplay between road networks and traffic dynamics, as drivers' routing choices significantly influence one another. City-related services, such as navigation services (e.g., TomTom) and mobility policies (e.g., road closures), impact traffic patterns and emissions. Navigation services can unintentionally increase emissions when many vehicles converge on the same routes, while mobility policies may have counterintuitive effects on traffic. We propose a simulation framework to assess the impact of road closure policies and navigation services on the urban environment. We use this framework and find that targeted road closures in Milan can reduce emissions by up to 10%, while others can

increase emissions by nearly 50%. Then, we examine navigation services' impact on vehicular traffic and CO₂ emissions, finding that they reduce emissions at low traffic loads. However, at high traffic loads and penetration rates, they cause conformist behavior, leading to inefficiencies and potentially higher emissions. To mitigate the conformist behavior induced by navigation services and reduce CO₂ emissions, we propose three solutions: (i) an individualistic approach using existing Alternative Routing (AR) algorithms, (ii) Metis, a coordinated solution that coordinates drivers and dynamically estimates traffic to diversify routes, and (iii) Polaris, an individual AR algorithm which considers road popularity to optimize traffic distribution. Motivated by the varying effectiveness of AR solutions across cities, we study cities' route diversification, defining

shortest path instability and introducing diverCity, a metric to assess a city's propensity towards route diversity. Analysis shows that diverCity benefits from extensive road networks, leading to less congestion. We also address the impact of mobility attractors on diverCity and propose mitigation strategies. This thesis comprehensively studies vehicular traffic dynamics, offering a simulation framework to evaluate the environmental impact of mobility policies and navigation services. In addition, it presents solutions to mitigate negative impacts and proposes metrics to quantify a city's potential to offer route diversity.

Heterogeneous Transfer Learning in Natural Language Processing

Author: Andrea Pedrotti, University of Pisa

Supervisors: Alejandro Moreo, Cnr-Isti, Fabrizio Sebastiani, Cnr-Isti

With the advances in Deep Learning, the term Transfer Learning (TL) has become ubiquitous in the field of Machine Learning. One of the most widely adopted strategies when working with pre-trained models is to fine-tune them on downstream tasks by leveraging a relatively smaller labeled dataset compared to the amount of training data used for the pre-training phase. Fine-tuning is in fact a common technique of transfer learning.

In general TL, refers to a set of techniques and approaches which leverage training data sampled from a source distribution to improve performance on a test set, the target, containing elements sampled from a different, but related, distribution. This paradigm brings about two major advantages. First, it increases performance on the target domain by making the algorithm more robust and resilient, allowing us to leverage powerful pre-trained models that are trained on hardware not widely available. Second, it allows the application of data-intensive techniques to many scarce-resource

domains where training an ad-hoc solution would be impossible.

In this thesis, we explore applications of Heterogeneous Transfer Learning (HTL) to the field of Natural Language Processing (NLP). We identify two main exploratory spaces: (i) the heterogeneous space defined by different languages (multilinguality) and (ii) the heterogeneous space defined by the intersection of languages and perceptual (multimodality) information. Lastly, we explore the benefits of HTL when dealing simultaneously with both multimodality and multilinguality.

Concerning (i), we address the task of Cross-Lingual Text Classification, where documents written in different languages must be classified according to a common classification scheme. In this scenario, we propose gFun, a two-tier architecture that leverages the space of posterior probabilities to represent documents written in different languages in a common language-agnostic space.


Regarding (ii), we focus on the recent convergence of language and perceptual modalities. In ViLMA, we assess the ability of current Video and Language Models (ViDLMs) to align textual representations with visual video-based ones. In our experiments, we find that ViDLMs do not perform better than Vision and Language Models (VLMs) working with images when presented with hard counterfactual examples. Through this benchmark, we underline the current limitations of ViDLMs and highlight to the research community central directions for future research.

Finally, we integrate both multilingual and multimodal aspects, focusing on the task of Multilingual Image-Text Classification (MITC), where items are associated with both textual and visual descriptions. Here, we extend the multilingual capabilities of gFun to the multimodal scenario and propose its variant called mm-gFun.


Welcome aboard!



Somnath Dutta
Research Staff (Tecnologo)
VC Lab



Andrea Pedrotti
Research Staff (Ricercatore)
AIMH Lab



Biagio Peccerillo
Research Staff (Tecnologo)
InfraScience Lab



Giovanni Puccetti
Research Staff (Ricercatore)
AIMH Lab


Conferences - Co-organized by ISTI



30th Annual ACM Conference on Intelligent User Interfaces (IUI 2025)
March 24-27, 2025, Cagliari, Italy
<https://iui.acm.org/2025/>



7th International Workshop on Big Mobility Data Analytics (BMDA 2025)
co-located with EDBT/ICDT 2025
March 25, 2025, Barcelona, Spain
<https://www.datastories.org/bmda25/index.html>



13th ACM/IEEE International Conference on Formal Methods
in Software Engineering (FormalISE 2025) co-located with ICSE 2025
April 27-28, 2025, Ottawa, Ontario, Canada
<https://conf.researchr.org/home/Formalise-2025>



6th ACM/IEEE International Conference on
Automation of Software Test (AST 2025) co-located with ICSE 2025
April 27-May 3, 2025, Ottawa, Ontario, Canada
<https://conf.researchr.org/home/ast-2025>



13th ACM/IEEE International Workshop on Software Engineering for
Systems-of-Systems and Software Ecosystems (SESoS 2025)
co-located with ICSE 2025
April 27-May 3, 2025, Ottawa, Ontario, Canada
<http://sesos2025.icmc.usp.br/>



13th International Workshop on Computational Intelligence for
Multimedia Understanding (IWCIM) co-located with IEEE ISCAS 2025
May 25-28, 2025, London, UK
<https://iwcim.itu.edu.tr/>



10th International Symposium on End-User Development (IS-EUD 2025)
June 16-18, 2025, Munich, Germany
<https://iseud2025.ubicomp.net/>



17th Advanced Infrared Technology and Applications (AITA) Conference
September 15-19, 2025, Kobe, Japan
<https://english.jsndi.jp/aita2025/index.html>



ISTI News is published by the Institute of Information Science and Technologies "A. Faedo"
National Research Council
Via G. Moruzzi 1-56124 Pisa Italy
<http://www.isti.cnr.it>

ISSN: 2724-6566
DOI: 10.32079/ISTINews

Editorial Secretariat

segreteria scientifica@isti.cnr.it

Editorial Board

Liana Falconetti
Silvia Giannini
Anna Molino
Claudia Raviolo
Roberto Scopigno
Maurice ter Beek
Adrian Wallwork

Layout and Design

Beatrice Rapisarda

Copyright notice

All authors, as identified in each article, retain copyright of their work.
The authors are responsible for the technical and scientific contents of their work

Privacy statement

The personal data (names, email addresses...) and the other information entered in ISTI News will be treated according with the provision set out in Legislative Degree 196/2003 (known as Privacy Code) and subsequently integration and amendment

Legal representative of the Institute of Information Science and Technologies "A. Faedo": Dott. Roberto Scopigno

ISTI News is not for sale but is distributed for purposes of study and research and published online at
<https://www.isti.cnr.it/en/isti-news>

To subscribe: send your email to istinews@isti.cnr.it