

ISTI News

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

ISSUE 13

June 2023

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

Cover Story
Making Open Science
the New 'Normality'

EDITORIAL
Conversing with Machines

Photo Credit Liana Falconetti (Isti-Cnr)

Contents

Editorial

Conversing with machines 3

Cover Story

Making Open Science the new ‘normality’ 6

New Projects

EFRA 9

Blue-Cloud 2026 10

TRANTOR 11

SciLake 13

CRAEFT 13

CREXDATA. 14

STARWARS. 15

Green.Dat.AI 15

ReefSurvAI 16

CRAFT-OA 17

EUCAIM 18

GraspOS. 19

AI-RIDE. 20

Selected Papers

Artificial intelligence of things at the edge: scalable and efficient distributed learning for massive scenarios. 21

Evaluating the impact of anchors deployment for an AoA-based indoor localization system 22

Can we communicate? Using dynamic logic to verify team automata. 23

A runtime environment for contract automata 24

Zero-shot learning for requirements classification: an exploratory study. 25

Minimisation of spatial models using barnching bisimilarity. 25

DevOpRET: continuous reliability testing in DevOps 26

Implicit reward structures for implicit reliability models. . . 26

From unstructured texts to semantic story maps 27

Measuring fairness under unawareness of sensitive attributes: A quantification-based approach. 27

Using Semantic Web to create and explore an index of toponyms cited in medieval geographical works 28

Generalized funnelling: ensemble learning and heterogeneous document embeddings for crosslingual text classification 29

Understanding user needs in smart homes and how to fulfil them 30

Global-scale parameters for ecological models. 30

A self-training automatic infant-cry detector 31

Estimating hidden fishing activity hotspots from vessel transmitted data 32

Feature-rich multiplex lexical networks reveal mental strategies of early language learning 33

A survey on deep learning for human mobility 34

An optimal algorithm for finding champions in tournament graphs 35

Distilled neural networks for efficient learning to rank. . . . 35

Raman spectroscopy and topological machine learning for cancer grading 36

Restoration and content analysis of ancient manuscripts via color space based segmentation 36

Shallow portion of an active geothermal system revealed by multidisciplinary studies: the case of Le Biancane (Larderello, Italy) 37

Efficient lung ultrasound classification 38

NoR-VDPNet++: real-time no-reference image quality metrics. 39

Texture inpainting for photogrammetric models 40

Dynamic behaviour of the carillon tower in Castel San Pietro Terme, Italy 41

Deep learning for structural health monitoring: an application to heritage structures. 41

The short-term effects of the Cosmos 1408 fragmentation on neighboring inhabited space stations and large constellations 42

Awards and Achievements

The VISIONE video retrieval system - runner-up at VBS 2023. 43

Thomas Alderighi awarded for the best Ph.D. thesis at University of Pisa 43

Gaia Pavoni winner of the VRVis Visual Computing award 2023 44

The AIMH Lab wins a DH award 2022. 44

Roberto Scopigno awarded the Eurographics gold medal 2023. 45

AST 2023 1st runner-up paper award. 45

IAASS Vladimir Syromyatnikov Safety-by-Design award 2023 46

Francesco Laccone is the winner of the 2023 Pier Luigi Nervi Prize 46

Miscellaneous

Ph.D. dissertations 47

Welcome aboard! 50

Conferences - Co-organized by ISTI 51

Conversing with machines

The rise of ChatGPT and the future of research on Large Language Models



The tension between public and private interests on recent deep learning technologies. "Amor sacro e amor profano" by Tiziano, redrawn by Stable Diffusion and ControlNet, using the prompt "Robots"

ChatGPT is a current buzzword in mainstream news. Thanks to their astonishing results, Large Language Models (LLMs) and the Generative Pretrained Transformer (GPT) have crossed the borders of AI labs to quickly become mass-adopted tools. How have we got to this? What's next?

From a nice spring to a torrid summer

In the editorial for the second issue of ISTI News, in 2017, Fabrizio Falchi and I wrote: "Deep Learning is not a blanket solution for all AI problems, but it is certainly a new approach to machine learning that will have a relevant impact on a wide range of applications. Deep Learning is here to stay". Seven years on, we can say we were right, and maybe even too cautious, at least in view of the events of these recent months.

2017 saw the full bloom of the AI spring, which, by convention, began in 2012 with the resounding success in an image classification challenge of the Convolutional Neural Networks of AlexNet. In recent years, research on deep learning has produced a continuous progression of new network

architectures, algorithms, and "tricks" (e.g., ResNet, Adam, ReLU...), which have shaped our knowledge on neural networks.

In tasks such as vision and image processing, deep learning has scored a sequence of impressive results, gaining the front spot in the news for the sensational image generation capabilities of diffusion-based methods, such as Stable Diffusion (SD, 2020).

Similarly, research on LLMs has evolved through a series of key innovations. The start of the LLM revolution can be identified with Google's Word2Vec (2013), a rather basic network that, by working on large volumes of text, was shown to be able to capture the semantics of words by learning to predict a hidden word from the words that appeared in its vicinity.

The next breakthrough was the transformer model (2017), with the idea of using the context of the occurrence of words also at prediction time, and not only at learning time. A number of LLMs based on the transformer model soon appeared, with GPT (2018) becoming the reference model for the generation of text.

The last step before ChatGPT was Instruct-GPT (2022), which added Reinforcement Learning (RL) to the recipe, making it possible to better drive the form and intent of the output.

The potential of LLMs had already become evident with GPT, but ChatGPT proved to be the killer application. ChatGPT (2023) found the "right" way to interact with humans: by chatting. After all, over the past decade text chats have become a dominant means of interaction between humans. Now the machine uses the same interface. A smart move!

What's next?

For the first time in decades, we are experiencing a valid alternative to the list of ten links we get from search engines. This could be a revolution in the way we represent and access knowledge, and it is not just about web search. GPT-based models can easily act as interfaces to databases, as automatic translators, as experts with encyclopedic knowledge, and as creative writers; this is a never seen before Swiss army knife of language and knowledge, with wide margins of

improvement and many functions yet to be discovered.

Unfortunately, models like GPT were already so large that most researchers are excluded from training them from scratch. Lacking the huge computing resources required, is it possible for the average researcher to have an impact on this field? Looking at the stream of papers published on the topic it could appear that the answer is yes, but there is an evident imbalance between the small number of entities (big IT companies, and the few academic institutions, typically collaborating with such companies) that can work on creating and studying the inner workings of LLMs, and the many that are limited to exploring the outputs of the available LLMs only by accessing them as (paid) black-box APIs.

Whether there is interest in discovering the full potential of LLMs, and their future evolution, or concern about the harm they can do, I do not believe that we should leave the exploration of this field to a few private (non-European) entities, thereby reducing ourselves to test-driving their outputs, or, even worse, to being their crash-test dummies.

There is an urgent need for dedicated funds and resources to support public research on deep learning. Enabling public research bodies to thoroughly investigate LLMs, and multimodal generative models in general, is our best chance to acquire fundamental knowl-

edge. We must shift the focus from the development of the next chat tool that attracts more users, to researching the fundamental properties and mechanisms of these large networks, and connecting this research to the way we humans represent and share our knowledge.

In my opinion, this is also the most promising approach to addressing the concerns about the potential damage these new technologies can do to our society. Let me comment that the harm from widespread disinformation predates these technologies, and is strongly linked to the spread of social networks. If news about Stable Diffusion and ChatGPT raises the awareness of the risks of disinformation, maybe we can ascribe this to them as a merit, and hope for a more informed and cautious use of these platforms as sources of reliable information.

With these goals in mind, what could the next hot research topics be?

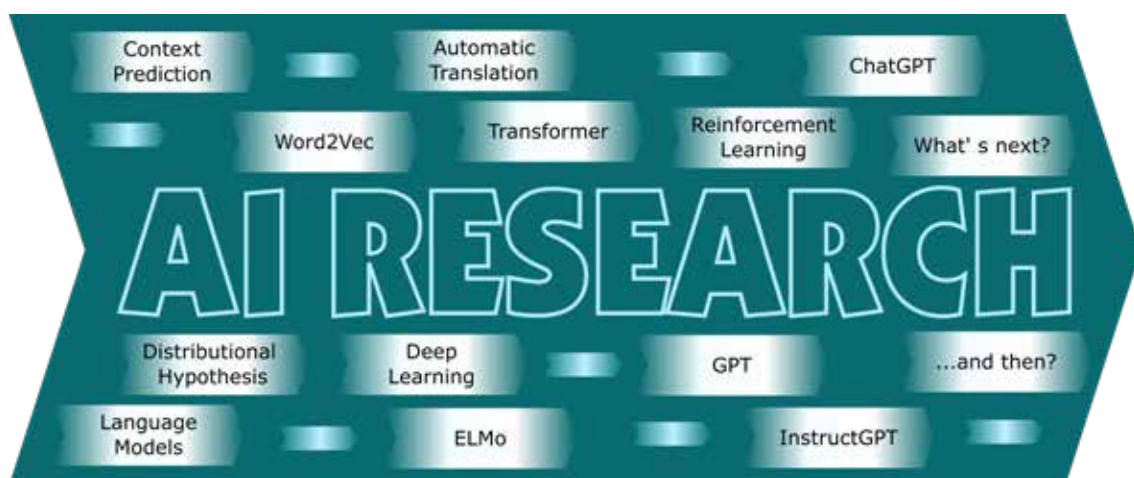
The timeline I outlined shows that ChatGPT is not a single algorithm that popped up from nowhere. The process that produces ChatGPT combines many learning paradigms: unsupervised learning, supervised learning, self-learning, and reinforcement learning. Many of these have been studied in isolation in the past. A holistic approach to all these paradigms could be the next big step in the field, especially if we want to solve the issues related to model hallucination and lack of

grounding (reporting as fact some information that is actually incorrect, false, or nonsensical).

The study of emerging behavior in larger LLMs, and the knowledge processing capabilities that show up when the size of the LLM is increased, may find interesting results when framed into complex system theory. This study may draw connections between what we observe in LLMs and phenomena already known from research in physics and biology. A special mention obviously goes to cognitive science, since the study of the nature and capabilities of the “intelligence” we feel by interacting with these models is likely to be more and more connected in the future to the study of human intelligence.

A more practical hot topic is the reduction of the computational cost involved in LLM-based training and inference, which would bring LLMs into the hands of more people. Another promising direction of research is the fusion of vision and language models into natively cross-modal models.

These are topics in fundamental research. In terms of applied research, we may soon experience a “gold rush” moment. Many of the disciplines that benefitted in the past from predictive algorithms may well obtain improvements simply by adopting the new learning models based on GPT (and RL, SD...).



ChatGPT is a major achievement, among many others, in the stream of AI research. What's next?

There could be many new opportunities for multidisciplinary collaborations.

One potential collaboration of importance is with law. Deep learning-based technologies are the target of criticism for alleged lack of ethics or dangers to society. Lawmakers feel pressured by public opinion to act on the topic. We must contribute by scientifically investigating these aspects, separating real issues from gratuitous fear mongering, and

indicating how to better manage the former.

Last but not least, we must also be science communicators. ChatGPT is the most evident result of a fruitful season of deep learning research. The sudden and direct impact it has on our society can backfire if not supported by sufficient information on the purposes, capabilities, and (especially) the limits of these technologies. We risk losing the benefits and advantages that these

innovations can bring to society because of an overreaction to the fears and risks they pose, true or only imagined. The researchers' way to counter fears and risks is through knowledge and education.

Contact: Andrea Esuli, AIMH Lab
andrea.esuli@isti.cnr.it

I thought ChatGPT should have its word, so I passed it my text and asked "Do you want to add anything to the text?"; this is the unedited output:

The text you provided already covers the main aspects of the topic quite well. However, if you would like me to add something, I can briefly mention the ongoing research efforts in developing more efficient and environmentally friendly LLM architectures. These efforts aim to reduce the carbon footprint associated with training and inference processes, considering the significant computational resources required for training large-scale models. This focus on sustainability aligns with the broader discussions on responsible AI development and addressing the environmental impact of AI technologies.

Furthermore, it is worth highlighting that the advancements in LLMs have also led to discussions around potential misuse and the need for robust safeguards to prevent the generation of malicious or harmful content. Researchers and policymakers are actively exploring techniques to mitigate risks and ensure the responsible deployment of LLMs in various domains.

By including these aspects, the text can provide a more comprehensive overview of the current landscape and ongoing efforts in the field of LLMs.

Making Open Science the new ‘normality’

Open science means conducting scientific research collaboratively and transparently. The timely implementation of this paradigm has become mandatory to tackle big societal research challenges. Many policy initiatives at the European and international levels have thus been launched to support and incentivate it (UNESCO¹, G7², OECD³). In Europe, open science has received strong political support from the European Council which, since 2016, has addressed multiple aspects of its implementation in Communications to the European Parliament^{4,5,6}.

If open science is to become the norm, research processes and results must be available and sharable according to the principle “as open as possible, as closed as necessary”. In today’s context, in which science is primarily performed digitally, this requires managing the data, software, services and virtual laboratories used in the research process to allow others to access, validate, reproduce and reuse them.

Innovative IT infrastructures and services, and legal, sociological, and economic models must replace those currently in use to render open science both possible and sustainable.

In Europe, the most advanced enabler funded so far by the Commission to support this new scientific paradigm is the European Open Science Cloud (EOSC⁷), an ecosystem of infrastructures with services facilitating data sharing and reuse. The development of this instrument is a political priority of the European Commission and the Member States, who have included it as action one in the ERA Policy Agenda 2022-2024⁸.

ISTI has been largely involved in supporting the move toward an open approach to science since the first decade of this millennium, both at the international level through research and development aimed at new technological solutions and data infrastructures and at the Cnr level by actively participating in the definition of new open science-related policies, initiatives and tools.

The rest of this article briefly introduces these activities.

Changing the Scholarly Communication Model: the OpenAIRE Infrastructure

Scholarly communication practice is expected to change radically with the transition to the open science paradigm in a context where science is primarily digital. Articles are gradually losing their central role as research communication instruments. They are becoming one of the many complementary digital objects (data, software, processes, virtual labs, etc.) that should be made available to researchers and society to support understanding of how a scientific activity was carried out and how the results were produced. Scientific results are expected to be communicated “as early as possible” in the research process, “as openly as possible”, and “as FAIR^{9,10}, is possible”, not only within a discipline but also between disciplines and society at large.

The OpenAIRE infrastructure¹¹, whose design and technological development has been coordinated by ISTI, is a key enabler in this change in scholarly communication. Born as an infrastructure for monitoring the evolution of open access within projects funded by the EU Commission, today it is an effective component of EOSC. It provides access to a Scholarly Research Graph connecting 166 million publications, 59 million research data, and 202 thousand research software items from 13 thousand data sources. The infrastructure offers services that support access to this wealth of results, monitor them, and train on open science best practices.

However, the design, development and usage of the OpenAIRE large graph have raised multiple research challenges for the ISTI researchers who initially conceived it. Effective and efficient deduplication is a very difficult process due to the huge amount of research objects and sources in question^{12,13,14}. Chal-

lenges also relate to the discoverability of research outcomes¹⁵ and the seamless exchange of information between diverse initiatives regarding Scientific Knowledge Graphs¹⁶. Today, the OpenAIRE Research Graph is an open resource made available with the CC-by licence, and is exploited for different types of analysis. In particular, at ISTI, it is used in addressing research questions characterising the scientometrics and science of science fields and also to monitor different dimensions of Open Science such as data citation practices¹⁷, to credit allocation in supplementary material¹⁸, and to benchmark the usage of other scholarly communication platforms^{19,20}.

Collaboration and Transparency: the D4Science Infrastructure

Collaboration and transparency, at the core of the open science paradigm, demand the availability and sharing of the objects used and produced during a scientific workflow. In digital science, these requirements imply managing all research results according to practices that guarantee their FAIRness. Ensuring this property is usually complex and time-consuming for researchers since it entails implementing appropriate protocols and adding adequate contextual information. This process is only feasible and sustainable if the contextual information is collected automatically as the scientific workflow progresses.

The D4Science infrastructure^{21,22} developed and operated at ISTI, addresses collaboration and transparency as one of its key objectives. D4Science offers Virtual Research Environments (VREs) which support digital science workflows covering activities ranging from data collection to data preparation, analysis and publishing. Each VRE is a collaborative environment in which a team of scientists can perform their activities by sharing data, services, methods, etc. Although D4Science was conceived more than fifteen years ago, it offers functionalities that intrin-

sically match today's open science practices. It natively supports "FAIRness-by-design" since contextual information is collected in the background as the scientific workflows progress. It also allows the sharing of outcomes progressively produced during the scientific process at different levels, from sharing among team members to sharing with the members of a review panel up until the full release of a product. D4Science currently operates 190+ VREs and 20 thematic gateways serving 21K+ users. It operates as the foundation for other domain-specific EU and research infrastructures like So-BigData²³ (social mining), and BlueCloud²⁴ (marine science). It also operates VREs serving other infrastructure initiatives, like the PNRR-funded ITINERIS²⁵ and FOSSR²⁶, organisations like FAO and IOTC, and ministerial initiatives like D4GNA²⁷.



The implementation of D4Science has necessitated facing research challenges at different levels of the infrastructure service stack. The design and development of integration patterns and solutions implementing the system-of-systems approach that is necessary to support the on-demand creation of VREs has been among the core problematics²⁸. Other challenges that are still at the centre of current research activities regard the development of co-creation-oriented facilitators enabling scientists to onboard their assets and share them easily with co-workers, the design and development of innovative and open-science-oriented scientific communication and sharing strategies, and the smooth and continuous integration with the scholarly communication graph developed by OpenAIRE.

A growing number of CNR colleagues from different institutes are now activating VREs to serve their research needs while meeting open science principles²⁹. Some of these VREs have been created to empower the activities of ISTI researchers. Notable examples are those used to produce open resources in the marine domain³⁰.

Open Science's Other Priorities

ISTI has not only contributed to the transition towards open science with leading technology supporting the open science paradigm, but it has also been visionary with regard to policies supporting open access to scientific results. ISTI was the first CNR institute to issue an open science policy and, consistently with this policy, to make available the scientific production of its researchers through a portal, the ISTI Open Portal^{32,32}. The solution implemented at ISTI has been adopted by two other CNR institutes (INO and ISPC), while others are in the pipeline. ISTI representatives have also contributed to shaping the CNR Open Science Roadmap³³, officially approved in April 2023. With this step, CNR aligned itself with other large European research institutes like CNRS, France, and Helmholtz, Germany. This roadmap is a concrete contribution towards implementing the Italian National Plan for Open Science³⁴.

Education and upskilling in Open Science are other areas where ISTI has been a pioneer by organising and teaching courses in universities and research institutes and contributing to the launch of the ICDI Open Science Competence Center. ISTI currently collaborates with other European research centres in the context of the EU project Skills4EOSC³⁵ dedicated to creating a network of competence centres to speed up the training of European researchers and harmonise the training of new professional figures for scientific data management. The plan for the future is to continue this path, encouraged by general agreement on its importance for the uptake of Open Science. The interest of ISTI researchers is centred in particular on the training of FAIR data management and

the corresponding tools. With this in mind, an analysis of ongoing tools and indicators is now underway to identify gaps and validate their practical applicability³⁶.

Since 2021 ISTI also operates the open-science.it portal, the point of reference for the national community on open science. The content of the portal is provided by the national community and its quality is assured by an editorial board including representatives of other national universities and research institutions.

Looking Ahead

We are well aware that implementing open science is a revolution that requires a profound change of mentality, policies and practices. We must admit that the understanding of this paradigm has still to be wholly shaped and that its full implementation will also depend on factors far beyond the scientific context, e.g. geopolitical and economic. What is clear is that its full adoption will require research in many areas, including juridical, sociological, economic and, primarily, informatics. ISTI, through its multiple and solid expertise in information technology and data management, is well placed to pursue this research successfully. This activity will necessarily be "translational", that is based on a bi-direction interplay between research and delivery and deployment processes that upgrade the research³⁷. It is this approach to research that will guarantee the necessary concrete acceleration required for open science to become the norm across science and society.

Contacts: Donatella Castelli,
InfraScience Lab
donatella.castelli@isti.cnr.it

References

1. "UNESCO Recommendation on Open Science", <https://www.unesco.org/en/open-science>
2. https://www8.cao.go.jp/cstp/kokusaiteki/g7_2023/annex1_os.pdf
3. Recommendation of the OECD Council concerning Access to Research Data from Public Funding, <https://www.oecd.org/science/inno/recommendation-access-to-research-data-from-public-funding.htm>
4. <https://data.consilium.europa.eu/doc/document/ST-9526-2016-INIT/en/pdf>
5. <https://data.consilium.europa.eu/doc/document/ST-13701-2021-INIT/en/pdf>
6. <https://www.consilium.europa.eu/media/56958/st10126-en22.pdf>
7. https://research-and-innovation.ec.europa.eu/strategy/strategy-2020-2024/our-digital-future/open-science/european-open-science-cloud-eosc_en
8. https://research-and-innovation.ec.europa.eu/system/files/2021-11/ec_rtd_era-policy-agenda-2021.pdf
9. FAIR acronym stands for "Findable", "Accessible", "Interoperable" and "Reusable",
10. Wilkinson, M., Dumontier, M., Aalbersberg, I. et al. The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* 3, 160018 (2016). <https://doi.org/10.1038/sdata.2016.18>
11. www.openaire.eu
12. Artini, Michele, De Bonis, Michele, Manghi, Paolo, Atzori, Claudio, Bardi, Alessia, & Baglioni, Miriam. (2021). OpenAIRE OpenOrgs Database (3.1.8). Zenodo. <https://doi.org/10.5281/zenodo.4751149>
13. De Bonis M, Manghi P, Atzori C. 2022. FDup: a framework for general-purpose and efficient entity deduplication of record collections. *PeerJ Computer Science* 8:e1058 <https://doi.org/10.7717/peerj-cs.1058>
14. <https://ceur-ws.org/Vol-3365/paper2.pdf>
15. <https://www.go-fair.org/implementation-networks/overview/discovery/>
16. <https://www.rd-alliance.org/groups/scientific-knowledge-graphs-interoperability-framework-skg-if-wg>
17. Ornella Irrera, Andrea Mannocci, Paolo Manghi, and Gianmaria Silvello. 2023. A Novel Curated Scholarly Graph Connecting Textual and Data Publications. *J. Data and Information Quality* Just Accepted (May 2023). <https://doi.org/10.1145/3597310>
18. Mannocci, Andrea, Irrera, Ornella, & Manghi, Paolo. (2022, September 7). Open Science and Authorship of Supplementary Material. Evidence from a Research Community. 26th International Conference on Science, Technology and Innovation Indicators (STI 2022), Granada, Spain. <https://doi.org/10.5281/zenodo.6975411>
19. Baglioni, M., Manghi, P., Mannocci, A. and Bardi, A., 2021. We Can Make a Better Use of ORCID: Five Observed Misapplications. *Data Science Journal*, 20(1), p.38. DOI: <https://doi.org/10.5334/dsj-2021-038>
20. Mannocci, A., Baglioni, M., Manghi, P. (2022). "Knock Knock! Who's There?" A Study on Scholarly Repositories' Availability. In: , et al. *Linking Theory and Practice of Digital Libraries. TPDFL 2022. Lecture Notes in Computer Science*, vol 13541. Springer, Cham. https://doi.org/10.1007/978-3-031-16802-4_26
21. www.d4science.org
22. M. Assante et al. (2019) Enacting open science by D4Science. *Future Generation Computer Systems*, Volume 101, Pages 555-563, doi:10.1016/j.future.2019.05.063
23. <http://www.sobigdata.eu/>
24. <https://blue-cloud.org/blue-cloud-2026>
25. <https://itineris.cnr.it/>
26. <http://www.fossr.eu/>
27. <https://gna.d4science.org/home>
28. M. Assante et al. (2022) Virtual research environments co-creation: The D4Science experience. *Concurrency Computat Pract Exper*. 2022:e6925. <https://doi.org/10.1002/cpe.6925>
29. <https://services.d4science.org/explore>
30. Coro, G., Bove, P. & Kesner-Reyes, K. Global-scale parameters for ecological models. *Sci Data* 10, 7 (2023). <https://doi.org/10.1038/s41597-022-01904-3>
31. <https://openportal.isti.cnr.it/>
32. M. Artini et al. (2020) RepOSGate: Open Science Gateways for Institutional Repositories. In: Ceci, M., Ferilli, S., Poggi, A. (eds) *Digital Libraries: The Era of Big Data and Data Science. IRCDL 2020. Communications in Computer and Information Science*, vol 1177. Springer, Cham. doi:10.1007/978-3-030-39905-4_15
33. Donatella Castelli, Giovanni De Simone, Flavia Cancedda, Leonardo Candela, Valentina Colcelli, Raffaele Conte, Francesca Di Donato, Silvia Giannini, Emma Lazzeri, Silvana Mangiaracina, Roberto Puccinelli, & Maria Adelaide Ranchino. (2023). Roadmap Scienza Aperta. https://doi.org/10.57665/BICE_ROADMAP2023
34. https://www.mur.gov.it/sites/default/files/2022-06/Piano_Nazionale_per_la_Scienza_Aperta.pdf
35. <https://www.skills4eosc.eu/>
36. D.Mangione, L. Candela L. and D. Castelli D. "A taxonomy of tools and approaches for FAIRification", <https://doi.org/10.5281/zenodo.6037508>
37. D. Abramson and M. Parashar, "Translational Research in Computer Science" in *Computer*, vol. 52, no. 09, pp. 16-23, 2019. doi: 10.1109/MC.2019.2925650

EFRA

Extreme Food Risk Analytics Funded by Horizon Europe

Through digitization and developments in sensor networks and Internet of Things (IoT) connectivity, the collection of data along the food supply chain has increased on an enormous scale. Not taking full advantage of this wealth of public and private data comes at a great cost. Despite best efforts and modern techniques, consumers world-wide may still get sick from foodborne diseases and food companies can consequently suffer huge economic and legal penalties from food product recalls. A new frontier in data-driven decision making has thus emerged: how to appropriately fuse public and private data, extract meaningful insights, and subsequently train Artificial Intelligence (AI) models to predict food safety risks before they affect human health. However, the huge volume and significant diversity and complexity of the data required to predict and prevent risks makes current technologies fail. The data resides in sources dispersed around the world, in many formats, types and languages, with missing values and insufficient context.

Extreme Food Risk Analytics (EFRA) aims to push these boundaries by exploring novel, experimental and promising approaches to extreme data mining, aggregation and

analytics technologies. EFRA aims to thoroughly test methods and tools that can significantly address the scientific, economic and societal challenges associated with the safety and quality of the food that European consumers eat. It will achieve this, through a number of key objectives:

1. Develop and test solutions able to discover and distil meaningful, reliable and useful food risk data from heterogeneous and dispersed/scarce data sources with minimal delay and appropriate format.
2. Design relevant interactions with users to measure usefulness for human risk prevention actions in real-world scenarios and use-cases.
3. Demonstrate solutions that enable the development of trustworthy, accurate, green and fair AI systems for food risk prevention.
4. Achieve ground breaking advances in performance and usefulness of food risk data discovery, collection, mining, filtering, and processing.
5. Integrate relevant technologies (e.g., big

data, IoT, High Performance Computing, AI) to help achieve its goals and foster links to respective communities of data innovators in the food supply chain.

6. Position its contributions within the overall ecosystem of public and private stakeholders that share data, technology and infrastructure to ensure the safety and quality of food in Europe.

To achieve these goals, EFRA will design, test, and deploy tools and undertake appropriate initiatives to facilitate their uptake, elicit feedback, and engage stakeholders in their design. The ultimate goal is to develop and deliver the integrated EFRA Platform to create the world's first analytics-enabled, secure-by-design, green data space for AI-enabled food risk prevention. The long-term aim is to enable the EU to become a global leader in the digital-led industry transition from reaction to food risk prevention.

Contact: Salvatore Trani, HPC Lab
salvatore.trani@isti.cnr.it
<https://efraproject.eu>



EFRA

EXTREME FOOD RISK ANALYTICS

Blue-Cloud 2026

A federated European FAIR and Open Research Ecosystem for oceans, seas, coastal and inland waters Funded by Horizon Europe

Over the past decades, Europe has developed an impressive capability for aquatic environmental observation, data-handling and sharing, modelling and forecasting, second to none in the world. This builds upon national environmental observation and monitoring networks and programs, complemented with EU infrastructures such as the Copernicus satellite observation programme and related thematic services, the European Marine Observation and Data Network (EMODnet), plus a range of environmental European Research Infrastructures and major R&D projects.

Within this framework, since October 2019, the pilot [Blue-Cloud project](#) combines the interests of the European Open Science Cloud (EOSC), which provides a virtual environment with open and seamless access to services for storage, management, analysis and re-use of research data, across borders and disciplines, and the marine research communities by developing a collaborative web-based environment based on the [D4Science](#) infrastructure. It provides sim-

plified access to an unprecedented wealth of multi-disciplinary datasets from observations, analytical services, and computing facilities essential for blue science.

Blue-Cloud 2026 aims at a further evolution of this pilot ecosystem into a Federated European Ecosystem to deliver FAIR & Open data and analytical services, instrumental for in-depth research of oceans, EU seas, coastal & inland waters. It develops a thematic marine extension to EOSC for open web-based science, serving the needs of the EU Blue Economy, Marine Environment and Marine Knowledge agendas.

Over the course of 42 months starting in January 2023, Blue-Cloud 2026 will not only evolve the core services, integrating more blue analytical services, and configuring more Virtual Labs. Rather, it will improve services for uptake of new data sets from a multitude of data originators (such as SeaDataNet, EurOBIS, Euro-Argo, ELIXIR-ENA, SOCAT, EcoTaxa, and ICOS-Ocean), and major e-infrastructures, namely EUDAT,

D4Science, and WEkEO (CMEMS DIAS) and for discovery and access to their structured data collections.

Blue-Cloud 2026's overall objective is to expand the federated approach of Blue-Cloud, involving more aquatic data stakeholders, and interacting with EOSC developments, in support of the EU Green Deal, UN SDG, EU Destination Earth, and the EU Mission Starfish on healthy oceans, seas, coastal and inland waters, ultimately to provide a core data service for the Digital Twin of the Ocean.

Blue-Cloud 2026 is co-coordinated by CNR, Trust-IT & MARIS, counting on a core team of partners such as VLIZ, Ifremer, MOI, Seascope Belgium. Overall it mobilises a solid, multidisciplinary, committed team of 40 partners from 13 EU countries.

Contact: Pasquale Pagano,
InfraScience Lab
pasquale.pagano@isti.cnr.it



TRANTOR

5G+ evoluTion to mutioRbitAI multibaNd neTwORks

Funded by Horizon Europe

European industry must invest in research, innovation for local businesses, to reduce reliance on third countries for satellite broadband internet access. Automating the satellite radio access network using Artificial Intelligence (AI) and Machine Learning (ML) can reduce manual processes, and unifying radio access and satellite operations control

is essential for efficient communications. TRANTOR envisages scenarios for Enhanced Mobile Broadband usage that provide connectivity to under-served areas and offer network resilience using terrestrial (TNs) and non-terrestrial networks (NTNs). TRANTOR will test a regenerative-satellite-based Next-Generation Radio Access Network ar-

chitecture to improve satellite capacity and reduce fronthaul power and network delay while maintaining Quality of Service (QoS). NTN can complement TNs and provide efficient and resilient connectivity where a TN is not available or reliable.

Scenario 1: End-to-End single-band connectivity with a single Geostationary (GEO) satellite



Use case 1: Disaster relief

NTNs are crucial in providing connectivity in areas where TNs are challenging to deploy or non-existent. In emergency situations, such

as after an earthquake, satellite connectivity is essential for coordinating rescue efforts and restoring 5G networks to assist rescue

teams in their work, providing radio access networks in areas where terrestrial communications are unavailable.

Scenario 2: End-to-End single-band connectivity with a single Low Earth Orbit (LEO) Satellite

Use case 2: Access to Vehicle-Mounted User Equipment




In an emergency situation in an area where the local network coverage is insufficient for real-time communication, public safety

NTNs can improve local and wide-area public safety networks by establishing direct communication between emergency responder

terminals, thus ensuring service continuity and general safety.

Scenario 3: End-to-End single band connectivity with centralized/distributed unit (CU/DU) split



Use case 3: CU/DU split study for the coming 3GPP Release 18


The 5G gNodeB (gNB), the node in a cellular network that provides connectivity between the user equipment and the core network, has been split into two parts, gNB-DU and

gNB-CU, to reduce required fronthaul capacity and delay, while maintaining QoS and network flexibility. The gNB-DU is on board the satellite and the gNB-CU is on the ground

segment. Users on Side B may use different network services, such as voice or video, with different latency tolerances.

Scenario 4: Multi-band transmission from a single GEO satellite

Use case 4: Wide area public safety




Emergency responders such as police, fire brigade, and medical personnel have been equipped with UEs and an external multi-band satellite antenna transmission capability,

which will allow them to communicate using messaging, voice, and video services even in adverse weather conditions on Side B. Depending on service availability and weather

conditions, they will employ either Ku or Ka band to communicate with Side A, where other emergency responders will coordinate to ensure continuity of service.

Scenario 5: Multi-satellite, multi-band transmission using two GEO satellites.

Use case 5: NTN-TN Service continuity mobile UE



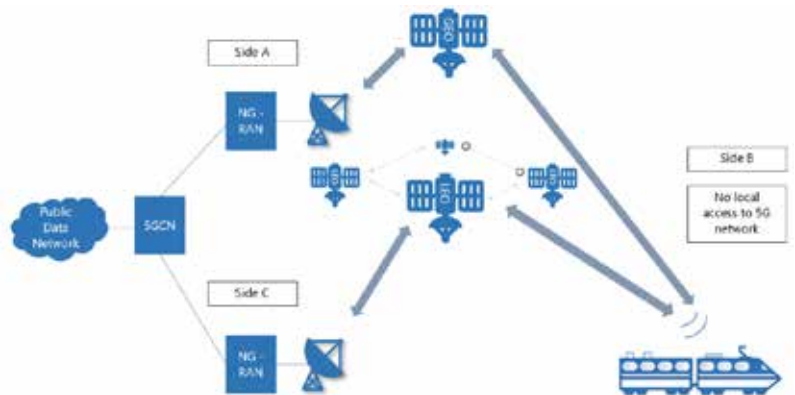
This use case focuses on the need for pervasive communications and ubiquitous coverage for connected vehicles, including ground

vehicles, maritime and aerial vehicles, and drones. These vehicles require continuous connectivity to support edge computing ser-

vices such as situation awareness, decision support and distributed AI tasks.

Scenario 6: Multiorbital, multi-band transmission using GEO and LEO satellites

Use case 6: NTN-TN Service continuity in trains



Long-distance trains often travel outside the coverage of terrestrial networks and rely on NTN such as GEO and LEO satellites for communication services. This combined NTN

use case allows for the provision of various 5G services to passengers, railway operators, and infrastructure operators.

Contact: Alberto Gotta, WN Lab
 alberto.gotta@isti.cnr.it
<https://www.trantor-he.eu/>

SciLake

Democratising and making sense out of heterogeneous scholarly content Funded by Horizon Europe

SciLake is a 3-year Horizon Europe research project that aims to extend technical work in the field of Science Knowledge Graphs (SKGs).

The project leverages SKGs as a basis to establish the concept of the scientific lake: a research ecosystem that facilitates the creating, combining, and querying of cross-domain and domain-specific SKGs. This ecosystem will include tools capable of extracting knowledge from unstructured (e.g., textual) information and tools for creating, interlinking, and managing SKGs. It will thus empower interoperability among SKGs, supporting knowledge transformation, unifying and simplifying how SKGs can be queried, and accelerating SKG graph processing and analysis.



The project will build a prototype and deliver an additional tier of smart services to assist the discovery of knowledge and improve research reproducibility. The work will be done in close consultation with four research communities (Neuroscience, Cancer, Transportation, and Energy). All services will be tailored to consider domain-specific requirements, ensuring their relevance to community needs and expectations. Each of these research communities will also demonstrate and evaluate selected SciLake services in the context of a domain-specific pilot. Finally, the project will leverage EOSC functionalities (e.g., comply with the Euro-

pean Open Science Cloud Interoperability Framework for monitoring, accounting, and Authentication and Authorization Infrastructures). It will integrate its open-source services into the OpenAIRE and EOSC Core services portfolio.

SciLake brings together a consortium of 13 partners from 9 different countries, including members with technological expertise in knowledge management and discovery to guarantee the successful implementation of the scientific lake and its services, and also domain experts from the four scientific disciplines that participate in the pilot activities

Contact: Miriam Baglioni, InfraScience Lab
miriam.baglioni@isti.cnr.it

CRAEFT

Craft Understanding, Education, Training, and Preservation for Posterity and Prosperity Funded by Horizon Europe

The Craeft project aims to advance our understanding of the various aspects of crafts as living and developing heritage, a sustainable source of income, and a means of expressing the mind through “imagery, technology, and sedimented knowledge”. Drawing on disciplines such as Anthropology, Knowledge Representation, Cognitive Science, Art History, Advanced Digitisation, Audiovisual & Haptic Immersivity, and Computational Intelligence, the project will take a generative approach that can accommodate digital conservation, reenactable preservation, and scaling of approaches for different materials and techniques.

Through the integration of digital technologies, such as intuitive aids, telecommunications, simulators, and high-end digitisation, Craeft will enhance craft education and training. This will help widen access,



economize learning, increase exercisability, and overcome the constraints of craft learning. The incorporation of haptics intelligence in digital design tools will help connect tacit knowledge with computer-aided craft-specific design. Workflow simulation will support experimental archaeology for the recovery of lost techniques, while analytic workflow analysis will provide digital fabrication opportunities for menial tasks, material savings and reuse, and reduction of energy consumption.

To enable certification and attachment of digital content, Craeft will introduce digital dimensions to individual craft works, which will also facilitate reputation and community building. A Community Portal will be established for professionals and the public, which will connect to the new media, host heterogeneous craft portfolios, and register practitioner credentials.

The Craeft project will be piloted in eight representative craft instances, with the objective of enhancing craft education and training, design, valorisation, and community services across a range of materials.

Contact: Valentina Bartalesi Lenzi,
AIMH Lab
valentina.bartalesi@isti.cnr.it
<https://www.craeft.eu>

CREXDATA

Critical Action Planning over Extreme-Scale Data Funded by Horizon Europe

The vision of CREXDATA is to develop a generic platform for real-time critical situation management including flexible action planning and agile decision making over data of extreme scale and complexity.

Emergency management and critical action planning call for timely and accurate decision making in several, diverse applications with the goal of optimizing economic, societal, and environmental impacts. Weather and health emergencies and maritime applications collect live streams of extreme data that can reach high speeds and volumes and possess highly volatile statistical properties, incorporating uncertainties and diverse modalities through a multitude of sources.

The CREXDATA project will develop the algorithmic apparatus, software architectures and tools for federated predictive analytics

and forecasting under uncertainty. The envisioned framework boosts proactive decision making providing highly accurate and transparent short- and long-term forecasts, explainable via advanced visual analytics and accurate, real-time, augmented reality facilities.

To achieve its vision, CREXDATA will develop a next generation Prediction-as-a-Service (PaaS) system where action planners can easily register their multimodal data stream sources and compute resource federations, and graphically design predictive and interpretable analytics workflows.

ISTI participates in the project contributing its expertise on explainable AI methods to make decision processes transparent, interpretable and accessible.

Three use cases will be used to evaluate project results:

- In the maritime domain, forecasting hazardous situations at sea.
- Weather emergency management, to allow authorities and first responders to engage proactively in the case of natural disasters.
- Health crisis management, to limit pandemic outbreaks and propose non-pharmaceutical means of patient treatment.

Contact: Salvatore Rinzivillo, KDD Lab
rinzivillo@isti.cnr.it
<https://crexdata.eu>

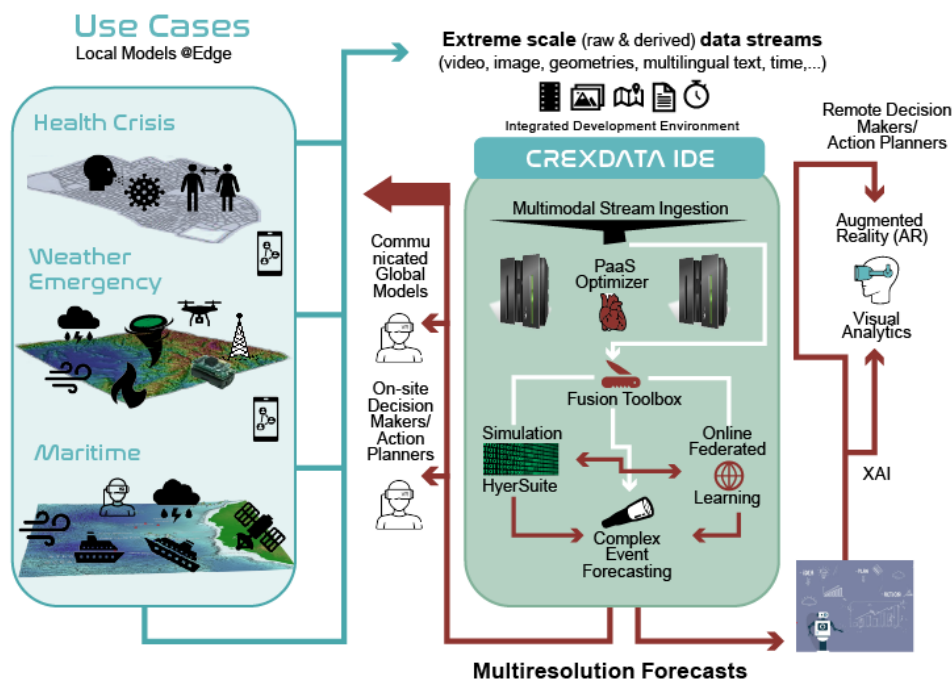


Figure 1: ©CREXDATA Concept Overview

STARWARS

STormwAteR and WastewAteR networkS heterogeneous data AI-driven management Funded by Horizon Europe, Marie Skłodowska-Curie Staff Exchange Action

Public and private stakeholders of the wastewater and stormwater sectors are increasingly faced with large quantities and multiple sources of information/data of different nature: databases of factual data, geographical data, various types of images, digital and analogue maps, intervention reports, incomplete and imprecise data (on locations and the geometric features of networks), evolving and conflicting data (from different eras and sources), etc.

Obtaining accurate and updated information on the underground wastewater and

stormwater networks is a challenge and a cumbersome task, especially in cities undergoing urban expansion.

Within this context, the main objective of the STARWARS project is to address the challenge by providing novel proposals for the management of heterogeneous data in stormwater and wastewater networks.

The project aims to bring together researchers from the AI and Water Sciences communities in order to enhance the emergence of new practical solutions for representing,

managing, modelling, merging, completing, reasoning, explaining and query answering over data of different forms pertaining to stormwater and wastewater networks.

The secondary objective is to produce new information and to promote knowledge sharing between the researchers involved in the project.

Contact: Umberto Straccia,
InfraScience Lab
umberto.straccia@isti.cnr.it

Green.Dat.AI

Energy-efficient AI-ready Data Spaces Funded by Horizon Europe

Green.Dat.AI aims to develop innovative energy-efficient large-scale data analytics services, ready to use in industrial AI-based systems. The project will demonstrate the efficiencies of these new analytics services in four industrial sectors: Smart Energy, Smart Agriculture/Agri-food, Smart Mobility, Smart Banking and six different application scenarios, leveraging the use of European Data Spaces.

The ambition is to exploit solutions already developed in recent H2020 projects and deliver an efficient, massively distributed, open-source, green, AI and federated learning-ready platform, and provide a validated go-to-market toolbox for AI-ready data spaces.

By design, the Green.Dat.AI Toolbox will be compliant with FAIR (Findable, Accessible, Interoperable and Reusable) data and



metadata management principles. In the long-term, the Green.Dat.AI platform will allow computing to move from data centers to edge devices, improving the accessibility of AI, and to shift computation from the cloud to personal devices, reducing the flow and potential leakage of sensitive data. It will thus enable the processing of data on the edge to eliminate transmission costs, leading to faster inference with a shorter reaction time and driving innovation in applications where these parameters are critical.

The Green.Dat.AI Consortium consists of a multidisciplinary group of 17 partners from 10 different countries (plus an associated party), and is well balanced in terms of expertise. Green.Dat.AI draws on knowledge from different domain experts in energy, transport, business and economics as well as Data Science and SW Engineering.

The KDD Lab of ISTI is contributing to Green.Dat.AI, working on a number of analytical tools and functionalities, including: explainable time series and mobility data predictive models, synthetic data generation, geographical transfer learning, and federated multi-objective learning

Contact: Mirco Nanni, KDD Lab
mirco.nanni@isti.cnr.it
www.greendatai.eu

ReefSurvAI

ReefSurvAI: towards a web-based AI infrastructure for coral reef surveying Funded by MAECI (joint international research project)

In January 2023, the Visual Computing Lab won a competitive call for a USA-Italy 2023-2025 international joint project, in the research area of Climate Change and resilience to Natural Disasters. The project is in collaboration with the Stuart Sandin Lab of the Scripps Institution of Oceanography of the University of California San Diego (UCSD). It regards the creation of a prototype AI-based web infrastructure to support coral reef investigations.


Nowadays, photogrammetry and Artificial Intelligence (AI) are fundamental in underwater investigations; high-resolution 3D image-based reconstructions have improved coral reef monitoring by facilitating novel seascape ecology analyses, and AI is accelerating image data interpretation, automatically counting, and measuring species of interest.

This rapid technological evolution has a major drawback; digital monitoring of coral reefs suffers from the lack of shared standards, data processing procedures, annotations, datasets to feed machine learning techniques and/or deep learning networks, and machine learning/deep learning models. This lack of replicability impacts the uniformity of analyses, complicating comparisons of scientific findings from different laboratories.

This joint project will study the feasibility of setting up an infrastructure dedicated to the digital investigation of coral reefs, including a repository for sharing datasets, resources, and machine learning/deep learning models. The introduction of such an infrastructure would greatly promote the adoption of common standards and tools for analysis, ensuring interoperability, scalability, and quality

assessment of image-based data. However, its design and implementation, fully implementing the FAIR principles of open science, poses several problems, ranging from scalability to homogeneity, data governance and access rights in a culturally and nationally broad community. In the three years of the project, the Visual Computing and the Stuart Sandin Labs will explore its feasibility by conducting pilot experiments to validate technologies and methodologies and providing prototypes that implement key portions of the infrastructure.

Contacts: G. Pavoni, M. Corsini, F. Ponchio, P. Cignoni, VC Lab
gaia.pavoni@isti.cnr.it
massimiliano.corsini@isti.cnr.it
federico.ponchio@isti.cnr.it
paolo.cignoni@isti.cnr.it



An example of coral taxonomy organized by color. The 3D model of the seafloor has been acquired through image-based techniques

CRAFT-OA

Creating a Robust Accessible Federated Technology for Open Access Funded by Horizon Europe

After several decades of evolution, Open Access (OA) publishing is now at the centre of scientific communication, providing access to scientific publications without barriers. In Diamond Open Access, authors can publish free of charge as the institutional sector (e.g. universities, research institutions, libraries) provide the necessary technological infrastructure. While the commercial model of Open Access dominates in anglophone journals from the Global North, the Diamond OA model shows a much higher level of diversity and origin. However, the Diamond OA landscape is fragmented, often underfunded, and not always sufficiently technically proficient to develop its full potential.

The CRAFT-OA project aims to respond to the European Commission's call to address the challenge of insufficiently tapped potential of institutional publishing and intends to contribute to the implementation of the Action Plan for Diamond Open Access presented by Science Europe, cOAlition S, OPERAS, and the French National Research Agency (ANR).

CRAFT-OA focuses on four threads of activities aimed at improving the technical and organisational infrastructure of Diamond OA:

- Provide technical improvements for journal platforms and journal software
- Build communities of practice to foster overall infrastructure improvement

- Increase visibility, discoverability and recognition for Diamond OA publishing
- Integrate Diamond OA publishing with EOSC and other large-scale data aggregators.

By offering tangible services and tools for the entire life cycle of journal publishing, CRAFT OA will empower local and regional platforms and service providers to upscale, become more professional and reach stronger interoperability with other scientific information systems for content. These developments will help organizations, researchers and editors involved in publishing to implement more sustainable and robust publishing workflows.

All CRAFT-OA's 23 consortium partners from 14 European countries are engaged in institutional publishing and relevant infrastructures, and are committed to sustaining and developing capacities in the field. Many

CRAFT-OA partners are leading national and European organisations in terms of open publishing and represent internationally visible centres of expertise in Open Science and FAIR principles (Findable, Accessible, Interoperable, Reusable), with strong connections to EOSC.

ISTI participates in this project as an affiliated party of OpenAIRE, a non-profit organisation founded in 2019 to maintain and operate a European infrastructure for Open Science, contributing the expertise of the InfraScience lab on interoperability of metadata and services, aggregative data infrastructures, integration with the European Open Science Cloud, and scientific knowledge graphs.

Contact: Alessia Bardi, InfraScience Lab
alessia.bardi@isti.cnr.it
<https://www.craft-oa.eu>

EUCAIM

Funded by Digital Europe

At the beginning of 2023, the European Commission launched the European Federation for Cancer Images (EUCAIM), a ground-breaking federated infrastructure deployment project aiming to power up imaging and AI towards precision medicine for Europe's cancer patients and citizens.

EUCAIM will address the fragmentation of existing cancer image repositories and establish a distributed Atlas of Cancer Imaging with over 60 million anonymised cancer image data from over 100,000 patients, accessible to clinicians, researchers and innovators across the EU for the development and benchmarking of trustworthy AI tools.

EUCAIM is the cornerstone of the European Commission initiated [European Cancer Imaging Initiative](#), a flagship of Europe's Beating Cancer Plan (EBCP), which aims to foster innovation and deployment of digital technologies in cancer treatment and care, to achieve more precise and faster clinical decision-making, diagnostics, treatments and predictive medicine for cancer patients. The initiative is led scientifically by Prof. Luis Martí-Bonmatí, Chairman of Radiology, La Fe University and Polytechnic Hospital (Valencia, Spain) and coordinated by the Euro-

pean Institute for Biomedical Imaging Research (EIBIR) in Vienna, Austria.

The project builds upon the results of the work of the "AI for Health Imaging" (AI4HI) Network which consists of 5 large EU-funded projects on big data and Artificial Intelligence in cancer imaging: Chaimeleon, EuCanImage, ProCancer-I, Incisive and Primage.

EUCAIM brings together 76 partners from 14 EU member states, covering competences in cancer imaging and care, big data in medical imaging, FAIR data management, ethical and legal aspects of medical data, development and deployment of research infrastructures, AI and machine learning

In line with the European data strategy and supporting the goals of the European Health Data Space, EUCAIM will partner with the AI Testing and Experimentation Facility for Health under the Digital Europe Programme.

CNR contributes to EUCAIM with two research institutes: the Institute of Biostructures and Biomedicine (IBB-CNR) based in Naples and Turin and the Institute of Information Science and Technologies "Alessandro Faedo" (ISTI-CNR) based in Pisa.

IBB-CNR participates as a core partner of Euro BioImaging, a European research infrastructure consortium, and will be leading the activities on "Business & Sustainability models" and also contributing to several tasks related to the central hub implementation and use cases. ISTI-CNR is contributing as a partner of the ProCancer-I project and is contributing to activities on the data federation & interoperability framework, by collaborating to the definition of a hyper-ontology and a common data model, leveraging MOLGENIS and OMOP. ISTI-CNR is also contributing to shaping the use cases for EUCAIM data and infrastructure, by defining user roles and user stories. The experience matured in ProCancer-I and the AI4HI initiative, with respect to AI model development and validation (<https://doi.org/10.48550/arXiv.2109.09658>) is being exploited to structure standardised processing pipelines for AI-based image processing.

Contact: Sara Colantonio, SI Lab
sara.colantonio@isti.cnr.it
<https://cancerimage.eu>



GraspOS

GraspOS: next Generation Research Assessment to Promote Open Science
Funded by Horizon Europe



The current research assessment system is often criticised for its over-reliance on quantitative metrics such as the Journal Impact Factor and h-index, which can incentivise researchers to prioritise publishing in high-impact journals over other vital aspects of scientific research, such as reproducibility, data sharing and other Open Science practices.

GraspOS is a 3-year Horizon Europe research project that seeks to address these issues with the ambitious goal of developing, assessing and putting into operation an open and trusted federated infrastructure for next-generation research metrics and indicators. It will offer data, tools, services and guidance to support and enable policy reforms for Open-Science-aware, responsible research assessment at researcher (individual/group), institutional, organisational and nation levels.

The project brings together a multidisciplinary consortium of 18 partners distributed over 10 European nations, including universities, research institutions, and non-profit organisations, with expertise in various areas such as data management, scholarly communication, open-access publishing, and research evaluation metrics development.

GraspOS results will include the following:

- The Open Science Assessment Framework (OSAF), a living and collaborative guide for indicator toolboxes capturing Open Science practices and other valuable factors for responsible research assessment.
- A set of Openness Profiles, which are fit-for-purpose OSAF-based templates for responsible research assessment, and an online database (OSAF Registry) that collects these profiles and case studies to promote experience sharing and mutual learning.
- A set of AI-driven metadata enrichment tools and services for enhancing missing attributes and novel indicator values on research outputs and enriching the links between them with semantics.
- A set of Open Science monitoring tools and services offering Open Science indicators, functions for manual annotations (corrections/edits/additions) and relevant evidence to support assessment processes.
- An EOSC-integrated Federated Open Metrics Infrastructure that extends the European Open Science Cloud Interoperability Framework with the appropriate metadata standards, protocols, components, and APIs to support AI-driven annotation/enrichment and to implement an EOSC Core Metrics service.
- A Community of Practice of responsible research assessment experts from relevant networks.
- A set of carefully designed training materials accessible via the EOSC Knowledge Hub.

A key aspect of GraspOS is its focus on community engagement. The project aims to bring together research assessment experts and related initiatives to co-develop the new framework. This will involve networking opportunities and sharing results and lessons learned from different initiatives. GraspOS recognises the fundamental role that communities of practice play in shaping the current research assessment system and in transitioning towards more responsible and more open assessment practices. The project will liaise with the Council for National Open Science Coordination (CoNOSC, <https://conosc.org>) and the OPUS (<https://opusproject.eu>) project and implement synergies with the Directorate-General for Research and Innovation (DG-RTD) coalition on reforming research assessment. The project will also establish liaisons with other relevant networks and Open Science communities, such as the network of OpenAIRE National Open Access Desks.

In summary, GraspOS is an ambitious project that seeks to transform how we assess scientific research by developing a more transparent, accountable, and effective Open-Science-aware assessment system in Europe. Its connection with EOSC and OpenAIRE will ensure the framework is widely adopted and integrated into the existing Open Science ecosystem.

Contact: Andrea Mannocci,
InfraScience Lab
andrea.mannocci@isti.cnr.it
<https://graspos.eu>
https://twitter.com/GraspOS_project

AI-RIDE

AI-RIDE: AI-driven Framework for Motorbike Driving Licence Exams Partially Funded by VEDLIoT European Project Open Call EU H2020 ICT-56-2020

AI-RIDE (Artificial Intelligence driven Riding Distributed Eye) is a multi-parameter AI-assisted telemetry system able to compute test scores and results for motorbike Practical Driving Courses (PDC) and Driving Licence Exams (DLE).

The AI multi-camera system recognizes most infractions that could jeopardize a driving license exam. The system uses computer vision logic which processes the video streams provided by the cameras. Monitored violations include coming into contact with a traffic cone, performing a non-compliant slalom between cones, going off-piste, placing a foot on the ground, failing the completion time constraints (there is a minimum limit for the short circuit and a maximum for the long one) and missing the correct final stop area.

The system's core consists of a deep learning (YOLOv5) custom model that recognizes four classes: Motorbike, Pilot, Cone, and XCone (cone on the ground), as depicted in the picture 1.

Object recognition shows a confidence of over 90% in varying weather and illumination conditions. Besides recognizing the various classes, after homography projection, the system can quantify the displacement of the same static element (e.g., a cone) in two consecutive frames: thus recognizing cone contact infraction. A Kalman filter improves position of the motorbike in every frame of the acquired video. The Kalman filter makes it possible to eliminate most errors and achieve reliable tracking. Based on the correct position and speed assessment, the system can identify infractions regarding non-compliant slaloms, off-piste driving, and incorrect final halting.

In order to achieve a good balance between overall costs and results, and in consideration of the speed of the vehicles in the fastest segments of the circuit, the system is based on cost-effective hardware: each node is made of IP POE camera FullHD@50fps with a dedicated GPU NVidia A4000. This combination makes it possible to process a full HD video of 2K frames in less than 2 minutes.

In addition to the primary infraction recognition, the system aims to be helpful in the training/learning phase of driving courses. It suggests specific improvements based on an analysis of the trajectories taken by the learner, classifying driving behaviour, and creating a large dataset of key performance parameters for driving test classification. For this purpose, as a first step, the path taken by the motorbike and all related data are represented in a virtual 3D world (see the picture). The data recorded is used for virtual world construction; the wealth of detail provides the instructors with a 3D digital map to analyze a student's test drive and explain and correct their errors (fig. 2).

It is hoped that the results of this research will enable the European Union to set a new standard for the European Driving Licence Examination for motorbikes, with a system that provides a fair, neutral, and homogeneous evaluation in all cases.

**Contacts: Marco Righi,
Giuseppe Riccardo Leone, SI Lab**
marco.righi@isti.cnr.it
giuseppe.leone@isti.cnr.it



Fig. 1



Fig. 2

Artificial intelligence of things at the edge: scalable and efficient distributed learning for massive scenarios

S. Bano, N. Tonello, P. Cassarà, A. Gotta

Computer Communications, vol. 205. Elsevier, 2023.

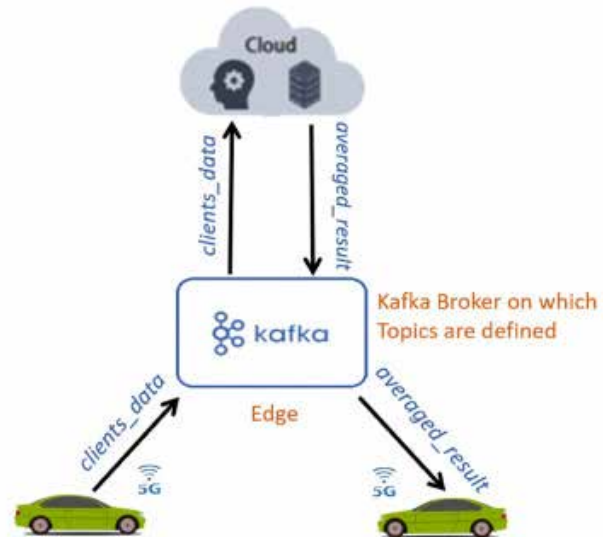
Federated Learning (FL) is a distributed optimization method in which multiple client nodes collaborate to train a machine learning model without sharing data with a central server. However, communication between numerous clients and the central aggregation server to share model parameters can cause several problems, including latency and network congestion. To address these issues, we propose a scalable communication infrastructure based on Information-Centric Networking built and tested on Apache Kafka®. The proposed architecture consists of a two-tier communication model. In the first layer, client updates are cached at the edge between clients and the server, while in the second layer, the server computes global model updates by aggregating the cached models. The data stored in the intermediate nodes at the edge enables reliable and effective data transmission and solves the problem of intermittent connectivity of mobile nodes. While many local model updates provided by clients can result in a more accurate global model in FL, they can also result in massive data traffic that negatively impacts congestion at the edge. For this reason, we couple a client selection procedure based on a congestion control mechanism at the edge for the given architecture of FL. The pro-

posed algorithm selects a subset of clients based on their resources through a time-based backoff system to account for the time-averaged accuracy of FL while limiting the traffic load. Experiments show that our proposed architecture has an improvement of over 40% over the network-centric based FL architecture, i.e., Flower. The architecture also provides scalability and reliability in the case of mobile nodes. It also improves


client resource utilization, avoids overflow, and ensures fairness in client selection. The experiments show that the proposed algorithm leads to the desired client selection patterns and is adaptable to changing network environments.

DOI: 10.1016/j.comcom.2023.04.010

MAS fetches the model parameters of the client from the broker to perform an average and sends them back to the broker



Overview of the proposed ICN-based FL communication architecture.



Overview of the client selection mechanism in FL with the control agent at the edge.

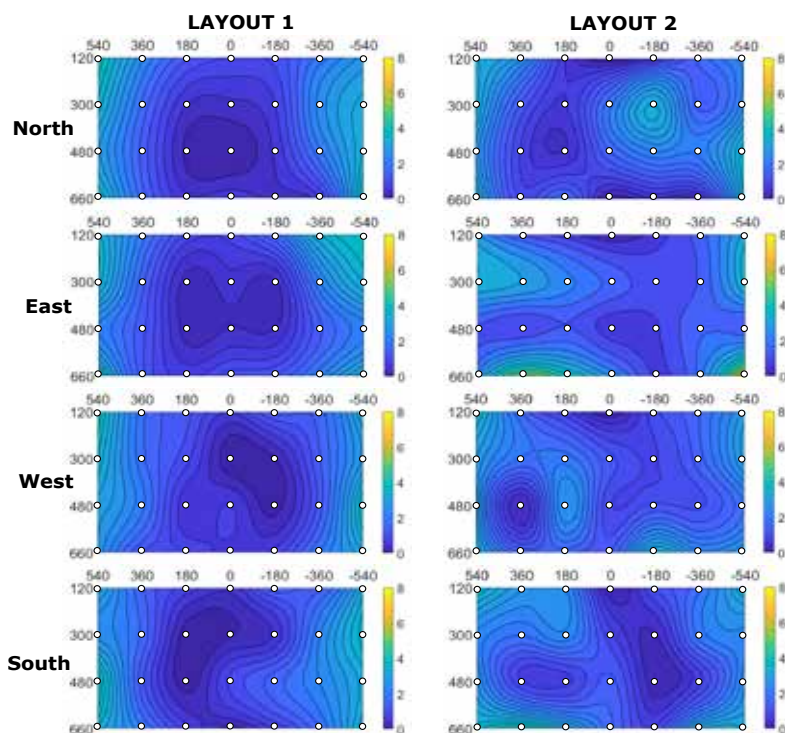
Evaluating the impact of anchors deployment for an AoA-based indoor localization system

F. Mavilia, P. Barsocchi, F. Furfari, M. Girolami

18th Wireless On-Demand Network Systems and Services Conference. IEEE, 2023.



The XPLR-AOA kit with anchor node and tag.



Localization error p expressed in meters for each layout and for different orientations.

Indoor localization techniques are rapidly moving toward the combination of multiple source of information. Among these, RSS, Time of Flight (ToF), Angle of Arrival (AoA) and of Departure (AoD) represent effective solutions for indoor environments. In this work, we propose an on-going activity investigating the performance of an indoor localization system based on the AoA-Bluetooth 5.1 specification, namely Direction Finding. We evaluate the effect of two anchor deployments and we test our localization algorithm by varying the orientation of the target according to four postures: North, West, South and East. From our study, we observe that anchor nodes deployed on the ceiling provide the best performance in terms of localization error. We conclude this work with a discussion of two further lines of investigation potentially increasing the performance of AoA-based indoor localization systems.

DOI: 10.23919/wons57325.2023.10061949

Can we communicate? Using dynamic logic to verify team automata

M.H. ter Beek, G. Cledou, R. Hennicker, J. Proença

25th International Conference on Formal Methods. Lecture Notes in Computer Science, vol. 14000. Springer, 2023.

Team automata describe networks of automata with input and output actions, extended with synchronisation policies guiding how many interacting components can synchronise on a shared input/output action. Given such a team automaton, one can reason over communication properties such as receptiveness (sent messages must be received) and responsiveness (pending receives must be satisfied). Previous work focused on how to identify these communication properties. However, verifying automatically these properties is non-trivial,

as it may involve traversing networks of interacting automata with large state spaces. This paper investigates (1) how to characterise communication properties for team automata (and subsumed models) using test-free propositional dynamic logic, and (2) how to use this characterisation to verify communication properties by model checking. A prototypical tool supports the developed theory, using an encoding to interact with the mCRL2 toolset for model checking. It can be executed online at <https://github.com/arcalab/team-a> and the screenshot de-

picts some of the available widgets, namely input: team automaton (1); output: encoded mCRL2 model and evaluated formulae (2); output: result of formula with counterexample (in this case) or witness (3); visualisation: composed team automaton (4), visualisation: individual component automata (5).

DOI: 10.1007/978-3-031-27481-7_9

The screenshot displays five widgets from the team-a tool interface:

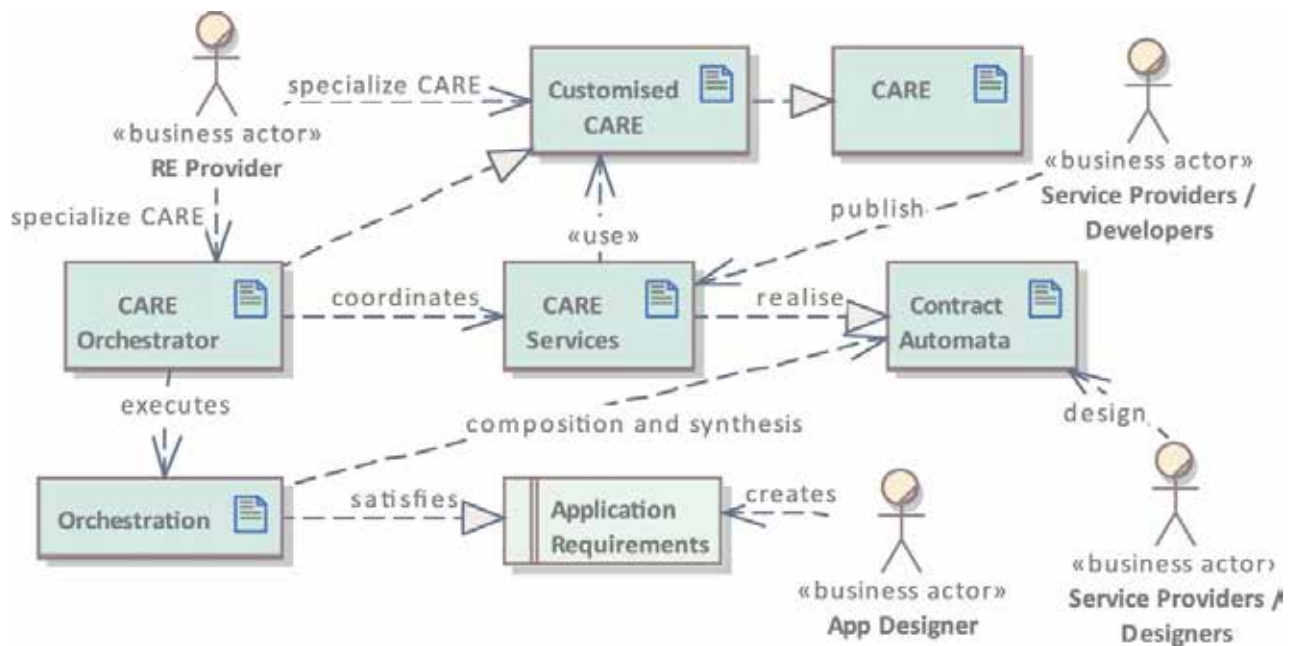
- Widget 1: ETA Specification** shows code for a race example with two runners and a controller.
- Widget 2: View mCRL2 evidence** displays the result "Responsiveness: false" and a state transition diagram for the mCRL2 model.
- Widget 3: Communication Properties' Characterisation in mCRL2** shows logical formulae for Receptiveness, Responsiveness, Weak Receptiveness, and Weak Responsiveness, along with the full mCRL2 system code.
- Widget 4: ETA diagram** shows a state transition diagram for the team automaton with states like (0,0,0) and (2,0,2).
- Widget 5: CA** shows two individual component automata: "runner" and "contr".

Screenshot of some of the widgets in the tool available online.

A runtime environment for contract automata

D. Basile, M.H. ter Beek

25th International Conference on Formal Methods. Lecture Notes in Computer Science, vol. 14000. Springer, 2023.



CARE methodology.

Contract automata have been introduced for specifying applications through behavioural contracts and for synthesising their orchestrations as finite state automata. This paper addresses the realisation of applications from contract automata specifications. We present CARE, a new runtime environment to coordinate services implementing contracts that guarantees the adherence of the implementation to its contract. We discuss how CARE can be adopted to realise contract-based applications, its formal guarantees, and we identify the responsibilities of the involved business actors. Experiments show the benefits of adopting CARE with respect to manual implementations.

Our proposal advances the state-of-the-art of the research on contract automata by showing a possible realisation of an orchestration engine, abstracted away in the

contract automata theory, but needed for implementing applications specified with contract automata, and guaranteeing that the implementation of service-based applications respect their specification. This contribution improves our understanding of the relation between a specification with contract automata and its implementation, and the corresponding level of abstraction.

With CARE, it is possible to promote a separation of concerns between formal methods experts specifying the expected behaviour using automata on one side, and developers (not required to be experts in formal methods) implementing the actions on the other. Furthermore, an application built using CARE is based on rigorous theoretical results from the contract automata theory, guaranteeing properties such as absence of deadlocks and absence of orphan messag-

es, reachability of final states, and absence of `ContractViolationException`. Moreover, CARE promotes modularity of applications composed by different services that are reusable in different applications and that can be adapted to satisfy different requirements through the synthesis of well-behaving orchestrations. Experiments showed the improvement in terms of decreased software complexity when using CARE instead of manually implementing the low-level interactions among services implementing the operations prescribed by their contracts.

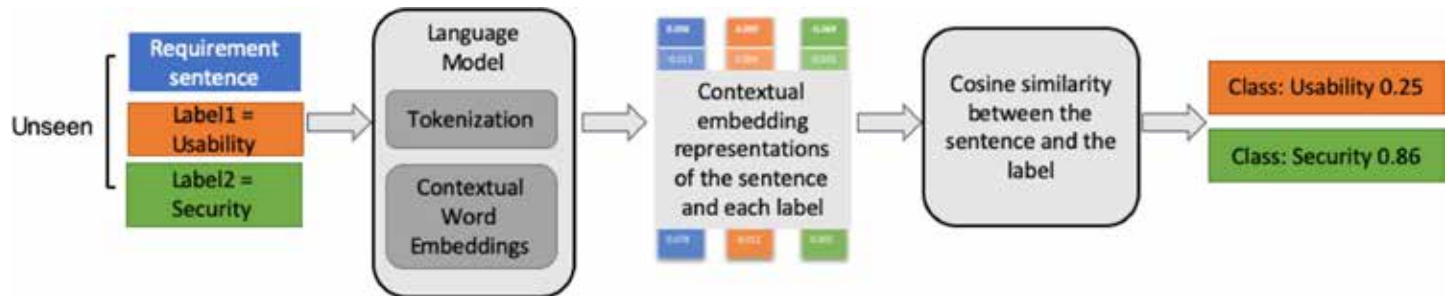
This publication received the available and reproducible badges from the artefact evaluation track.

DOI: 10.1007/978-3-031-27481-7_31

Zero-shot learning for requirements classification: an exploratory study

W. Alhoshan, A. Ferrari, L. Zhao

Information and Software Technology, vol. 159. Elsevier, 2023.



Zero-shot Learning for requirements sentences.

Requirements engineering (RE) researchers have been experimenting with machine learning (ML) and deep learning (DL) approaches for a range of RE tasks, such as requirements classification, requirements tracing, ambiguity detection, and modeling. However, most of today's ML/DL approaches are based on supervised learning techniques, meaning that they need to be trained using a large amount of task-specific labelled training data. This constraint poses an enormous challenge to RE researchers, as the lack of labelled data makes it difficult for them to fully exploit the ben-

efit of advanced ML/DL technologies. This paper addresses this problem by showing how a zero-shot learning (ZSL) approach can be used for requirements classification without using any labelled training data. We focus on the classification task because many RE tasks can be framed as classification problems. The ZSL approach used in our study employs contextual word-embeddings and transformer-based language models (LMs). We demonstrate this approach through a series of experiments to perform three classification tasks: (1) FR/NFR – classification functional requirements vs

non-functional requirements; (2) NFR – identification of NFR classes; (3) Security – classification of security vs non-security requirements. This study demonstrates the potential of ZSL for requirements classification. An important implication is that it is possible to have very little or no training data to perform classification tasks. The proposed approach thus contributes to the solution of the long-standing problem of data shortage in RE.

DOI: 10.1016/j.infsof.2023.107202

Minimisation of spatial models using branching bisimilarity

V. Ciancia, J.F. Groote, D. Latella, M. Massink, E.P. de Vink

25th International Conference on Formal Methods. Lecture Notes in Computer Science, vol. 14000. Springer, 2023.

Spatial logic and spatial model checking have great potential for traditional computer science domains and beyond. Reasoning about space involves two different conditional reachability modalities: a forward reachability, similar to that used in temporal logic, and a backward modality representing that a point can be reached from another point, under certain conditions. Since spatial mod-

els can be huge, suitable model minimisation techniques are crucial for efficient model checking. An effective minimisation method for the recent notion of spatial Compatible Path (CoPa)-bisimilarity is proposed, and shown to be correct. The core of our method is the encoding of Closure Models as Labelled Transition Systems, enabling minimisation algorithms for branching bisimulation

to compute CoPa equivalence classes. Initial validation via benchmark examples demonstrates a promising speed-up in model checking of spatial properties for models of realistic size.

DOI: 10.1007/978-3-031-27481-7_16.

DevOpRET: continuous reliability testing in DevOps

A. Bertolino, G. De Angelis, A. Guerriero, B. Miranda, R. Pietrantuono, S. Russo

Journal of Software: Evolution and Process, vol. 35. Wiley, 2023.

To enter the production stage, in DevOps practices candidate software releases have to pass quality gates, where they are assessed to meet established target values for key indicators of interest. We believe software reliability should be an important such indicator, as it greatly contributes to the end-user satisfaction. We propose DevOpRET, an approach for reliability testing as part of the acceptance testing stage in DevOps. DevOpRET relies on operational-profile-based testing, a common reliability



DevOps scope.


assessment technique. DevOpRET leverages usage and failure data monitored in operations to continuously refine its estimate. We evaluate accuracy and efficiency of DevOpRET through controlled experiments with a real-world open source platform and with a microservice architectures benchmark. The results show that DevOpRET provides accurate and efficient estimates of the true reliability over subsequent DevOps cycles.

DOI: 10.1002/smr.2298

Implicit reward structures for implicit reliability models

G. Masetti, L. Robol, S. Chiaradonna, F. Di Giandomenico

IEEE Transactions on Reliability, vol. 72. IEEE, 2023.



Classification of the models and reward structures on top of them according to explicit versus implicit representation and level of modeling abstraction. Newly introduced definitions for reward structures are underlined with a wavy line. The wavy arrow highlights one of the contributions of Section IV-A: how to project reward structures from high-level to low-level models.

such as stochastic Petri nets. Since the component models are implicit, the reward structure that characterizes the dependability properties has to be implicit as well. Therefore, we present a new formalism to specify those reward structures. The focus here is on component models that can be mapped to stochastic automata with one or several absorbing states so that the system model can be mapped to a stochastic automata network with one or several absorbing states. Correspondingly, the new reward structure defined on each component's model is mapped to a reward vector so that the dependability-related properties of the system are expressed through a newly introduced measure defined starting from those reward vectors. A simple, yet representative, case study is adopted to show the feasibility of the method.

A new methodology for effective definition and efficient evaluation of dependability-related properties is proposed. The analy-

sis targets the systems composed of a large number of components, each one modeled implicitly through high-level formalisms,

DOI: 10.1109/TR.2022.3190915

From unstructured texts to semantic story maps

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

International Journal of Digital Earth, vol. 16. Taylor&Francis, 2023.

Digital maps greatly support storytelling about territories, especially when enriched with data describing cultural, societal, and ecological aspects, conveying emotional messages that describe the territory as a whole. Story maps are interactive online digital narratives that can describe a territory beyond its map by enriching the map with text, pictures, videos, and other multimedia information. This paper presents a semi-automatic workflow to produce story maps from textual documents containing territory data. An expert first assembles one territory contextual document containing text and images. Then, automatic processes use natural language processing and Wikidata services to (i) extract key concepts (entities) and geospatial coordinates associated with the territory, (ii) assemble a logically-ordered sequence of enriched story-map events, and



The architectural schema of our workflow for semi-automatically creating story maps.

(iii) openly publish online story maps and an interoperable Linked Open Data semantic knowledge base for event exploration and inter-story correlation analyses. Our workflow uses an Open Science-oriented methodology to publish all processes and data. Through our workflow, we produced story maps for the value chains and territories

of 23 rural European areas of 16 countries. Through numerical evaluation, we demonstrated that territory experts considered the story maps effective in describing their territories, and appropriate for communicating with citizens and stakeholders.

DOI: 10.1080/17538947.2023.2168774

Measuring fairness under unawareness of sensitive attributes: a quantification-based approach

A. Fabris, A. Esuli, A. Moreo, F. Sebastiani

Journal of Artificial Intelligence Research, vol. 76. AI Access Foundation, 2023.

Algorithms and models are increasingly deployed to inform decisions about people, inevitably affecting their lives. As a consequence, those in charge of developing these models must carefully evaluate their impact on different groups of people and favour group fairness, that is, ensure that groups determined by sensitive demographic attributes, such as race or sex, are not treated unjustly. To achieve this goal, the availability (awareness) of these demographic attributes to those evaluating the impact of these models is fundamental. Unfortunately, collecting and storing these attributes is often in conflict with industry practices and legislation on data minimisation and privacy.

For this reason, it can be hard to measure the group fairness of trained models, even from within the companies developing them. In this work, we tackle the problem of measuring group fairness under unawareness of sensitive attributes, by using techniques from quantification, a supervised learning task concerned with directly providing group-level prevalence estimates (rather than individual-level class labels). We show that quantification approaches are particularly suited to tackle the fairness-under-unawareness problem, as they are robust to inevitable distribution shifts while at the same time decoupling the (desirable) objective of measuring group fairness from the

(undesirable) side effect of allowing the inference of sensitive attributes of individuals. More in detail, we show that fairness under unawareness can be cast as a quantification problem and solved with proven methods from the quantification literature. We show that these methods outperform previous approaches to measure demographic parity in five experimental protocols, corresponding to important challenges that complicate the estimation of classifier fairness under unawareness.

DOI: 10.1613/jair.1.14033

Using Semantic Web to create and explore an index of toponyms cited in medieval geographical works

V. Bartalesi, N. Pratelli, E. Lenzi, P. Pontari

Journal on Computing and Cultural Heritage, vol. 16. ACM, 2023.

Western thought in European history was mainly affected by the image of the world created during the Middle Ages and Renaissance. The most popular reason to travel during the Middle Ages was taking a pilgrimage. Jerusalem, Rome, and Santiago de Compostela were the most popular destinations. It is not surprising that a lot of works written by travellers as guides for pilgrims exist. By the beginning of the Renaissance, a more precise image of the world was defined thanks to the discovery of ancient geographical models, especially the work of Ptolemy. The three years (2020-2023) Italian National research project IMAGO - Index Medii Aevi Geographiae Operum - aims to provide a systematic overview of the medieval and renaissance Latin geographical literature using the Semantic Web technologies and the LOD paradigm. Indeed, until now, this literature has not been studied using digital methods. In particular, this paper presents how we formally represented the knowledge about the toponyms, or place names, in the IMAGO ontology. To maximise the interoperability, we developed the IMAGO ontology as an extension of two reference vocabularies: the CIDOC CRM and its extension FRBRoo, including its in-progress reformulation, LRMoo. Furthermore, we used Wikidata as reference knowledge base. As case study, we chose to represent the knowledge related to the toponyms cited by the Italian poet Dante Alighieri in his Latin works. We carried out a first experiment for visualising the knowledge about these toponyms on a map and in the form of tables and CSV files.

DOI: 10.1145/3582263



The toponyms reported by Dante Alighieri in his Latin works.

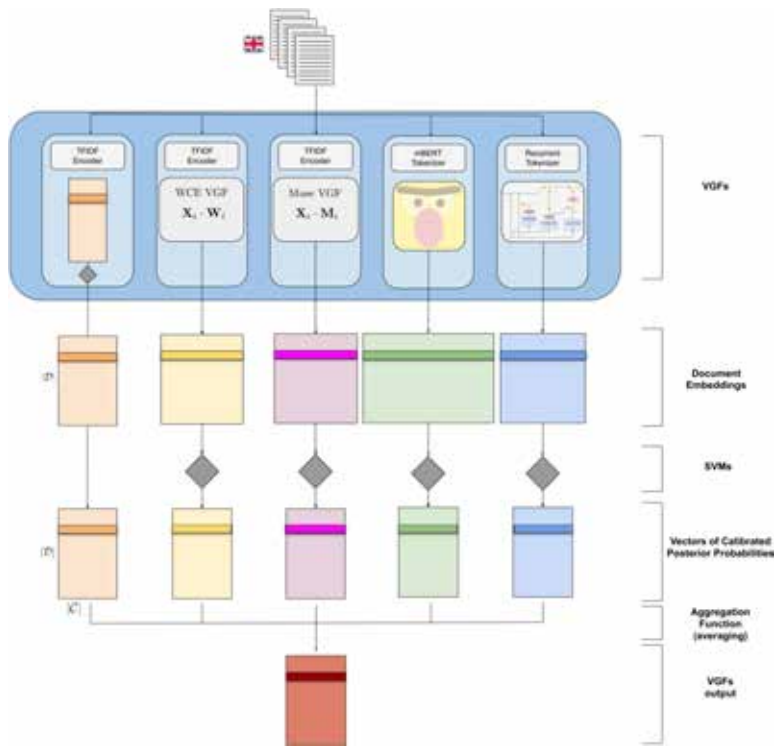


The data associated to Adria reported in a pop-up.

Generalized funnelling: ensemble learning and heterogeneous document embeddings for cross-lingual text classification

A. Moreo, A. Pedrotti, F. Sebastiani

ACM Transactions on Information Systems, vol. 41. ACM, 2023.



The averaging policy for view aggregation: The views are recast in terms of vectors of calibrated posterior probabilities before being averaged. Note that the resulting vectors lie in the same vector space. For ease of visualization, only one language (English) is shown.

Funnelling (Fun) is a recently proposed method for cross-lingual text classification (CLTC) based on a two-tier learning ensemble for heterogeneous transfer learning (HTL). In this ensemble method, 1st-tier classifiers, each working on a different and

language-dependent feature space, return a vector of calibrated posterior probabilities (with one dimension for each class) for each document, and the final classification decision is taken by a meta-classifier that uses this vector as its input. The meta-classifier

can thus exploit class-class correlations, and this (among other things) gives Fun an edge over CLTC systems in which these correlations cannot be brought to bear. In this paper we describe Generalized Funnelling (gFun), a generalisation of Fun consisting of an HTL architecture in which 1st-tier components can be arbitrary view-generating functions, i.e., language-dependent functions that each produce a language-independent representation (“view”) of the (monolingual) document. We describe an instance of gFun in which the meta-classifier receives as input a vector of calibrated posterior probabilities (as in Fun) aggregated to other embedded representations that embody other types of correlations, such as word-class correlations (as encoded by Word-Class Embeddings), word-word correlations (as encoded by Multilingual Unsupervised or Supervised Embeddings), and word-context correlations (as encoded by multilingual BERT). We show that this instance of gFun substantially improves over Fun and over state-of-the-art baselines, by reporting experimental results obtained on two large, standard datasets for multilingual multilabel text classification. Our code that implements gFun is publicly available.

DOI: 10.1016/j.ecoinf.2022.101675

Understanding user needs in smart homes and how to fulfil them

A. Mattioli, F. Paternò

9th International Symposium on End-User Development. Lecture Notes in Computer Science, vol. 13917. Springer, 2023.

Smart homes are becoming a widespread reality given the increasingly available number of connected objects and sensors. However, it is still unclear what people expect from automations that are made possible by this technological evolution. In addition, it is unclear whether current trigger-action pro-

gramming (TAP) languages offer sufficient operators and constructs to specify the desired automations. In this paper, we report on a study aiming to provide useful elements to address such issues. It involved 34 users without experience in IoT programming who created 204 desired home automations. We

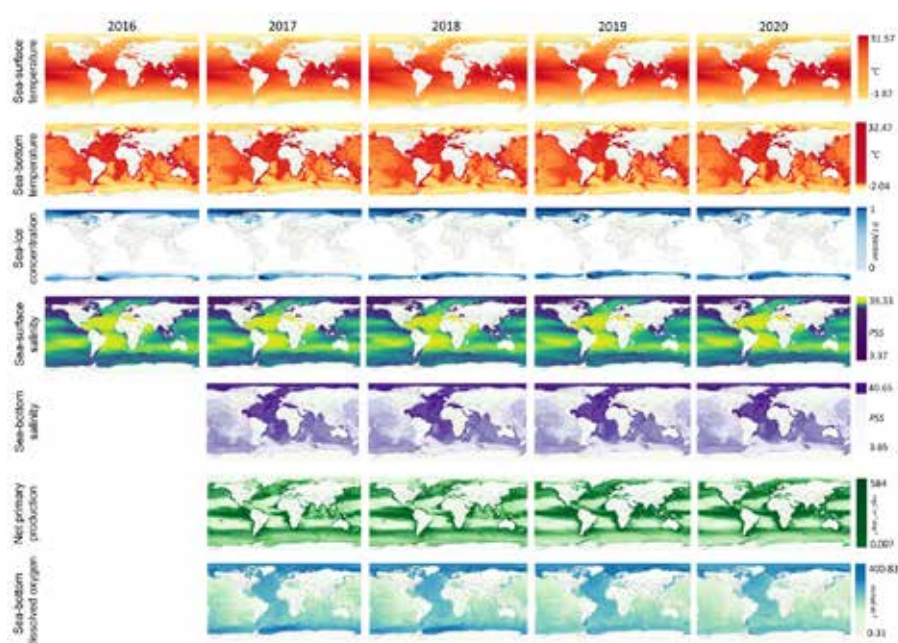
discuss an analysis of such results in terms of the relationships found between smart-home components and of the requirements for novel operators in TAP languages.

DOI: 10.1007/978-3-031-34433-6_8

Global-scale parameters for ecological models

G. Coro, P. Bove, K. Kesner-Reyes

Scientific Data, vol. 10. Springer Nature, 2023.



Comparison between the distributions of the environmental parameters used for time series and habitat analyses. The displayed maps have a global-scale 0.1° resolution.

This paper presents a collection of environmental, geophysical, and other marine-related data for marine ecological models and

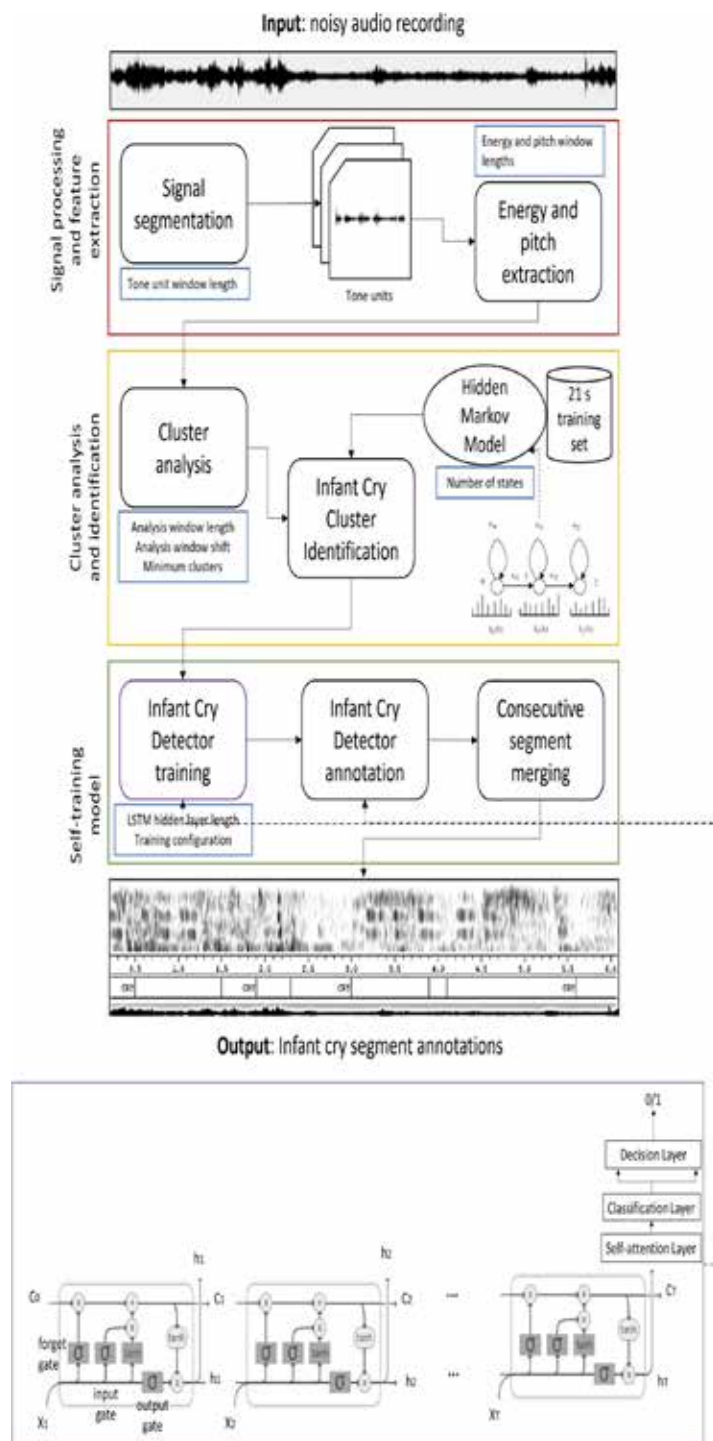
ecological-niche models. It consists of 2132 raster data for 58 distinct parameters at regional and global scales in the ESRI-GRID

ASCII format. Most data originally belonged to open data owned by the authors of this article but residing on heterogeneous repositories with different formats and resolutions. Other data were specifically created for the present publication. The collection includes 565 data with global scale range; 154 at 0.5° resolution and 411 at 0.1° resolution; 196 data with annual temporal aggregation over ~10 key years between 1950 and 2100; 369 data with monthly aggregation at 0.1° resolution from January 2017 to ~May 2021 continuously. Data were also cut out on 8 European marine regions. The collection also includes forecasts for different future scenarios such as the Representative Concentration Pathways 2.6 (63 data), 4.5 (162 data), and 8.5 (162 data), and the A2 scenario of the Intergovernmental Panel on Climate Change (180 data).

DOI: 10.1038/s41597-022-01904-3

A self-training automatic infant-cry detector

G. Coro, S. Bardelli, A. Cuttano, R.T. Scaramuzzo, M. Ciantelli
Neural Computing and Applications, vol. 35. Springer, 2023.



Schema of the proposed infant-cry detection workflow. The lowest frame shows the LSTM+A classification model.

Infant cry is one of the first distinctive and informative life signals observed after birth. Neonatologists and automatic as-

sistive systems can analyse infant cry to early-detect pathologies. These analyses extensively use reference expert-curated

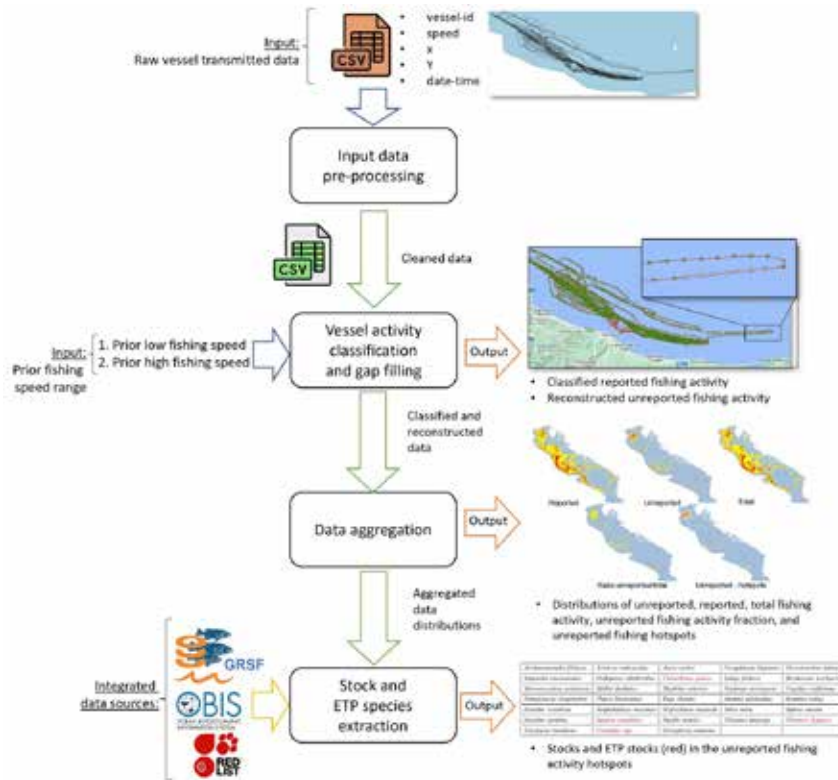
databases containing annotated infant-cry audio samples. However, these databases are not publicly accessible because of their sensitive data. Moreover, the recorded data can under-represent specific phenomena, or the operational conditions required by other medical teams. Additionally, building these databases requires significant investments that few hospitals can afford. This paper describes an open-source workflow for infant-cry detection, which identifies audio segments containing high-quality infant-cry samples with no other overlapping audio events (e.g., machine noise or adult speech). It requires minimal training because it trains an LSTM-with-self-attention model on infant-cry samples automatically detected from the recorded audio through cluster analysis and HMM classification. The audio signal processing uses energy and intonation acoustic features from 100-ms segments to improve spectral robustness to noise. The workflow annotates the input audio with intervals containing infant-cry samples suited for populating a database for neonatological and early diagnosis studies. On 16 min of hospital phone-audio recordings, it reached sufficient infant-cry detection accuracy in 3 neonatal care environments involving 20 infants subject to heterogeneous cry stimuli and had substantial agreement with an expert's annotation. Our workflow is a cost-effective solution, particularly suited for a sub-intensive care environment, scalable to monitor from one to many infants. It allows a hospital to build and populate an extensive high-quality infant-cry database with a minimal investment.

DOI: 10.1007/s00521-022-08129-w

Estimating hidden fishing activity hotspots from vessel transmitted data

G. Coro, L. Sana, C. Ferrà, P. Bove, G. Scarcella

Frontiers in Sustainable Food Systems, vol. 7. Frontiers, 2023.



Conceptual schema of our workflow.

Monitoring fishery activity is essential for resource planning and guaranteeing fisheries sustainability. Large fishing vessels constantly and continuously communicate their positions via Automatic Identification System (AIS) or Vessel Monitoring Systems (VMSs). These systems can use radio or Global Positioning System (GPS) devices to transmit data. Processing and integrating these big data with other fisheries data allows for exploring the relations between socio-economic and ecosystem assets in marine areas, which is fundamental in fishery monitoring. In this context, estimating actual

fishing activity from time series of AIS and VMS data would enhance the correct identification of fishing activity patterns and help assess regulations' effectiveness. However, these data might contain gaps because of technical issues such as limited coverage of the terrestrial receivers or saturated transmission bands. Other sources of data gaps are adverse meteorological conditions and voluntary switch-offs. Gaps may also include hidden (unreported) fishing activity whose quantification would improve actual fishing activity estimation. This paper presents a workflow for AIS/VMS big-data analysis

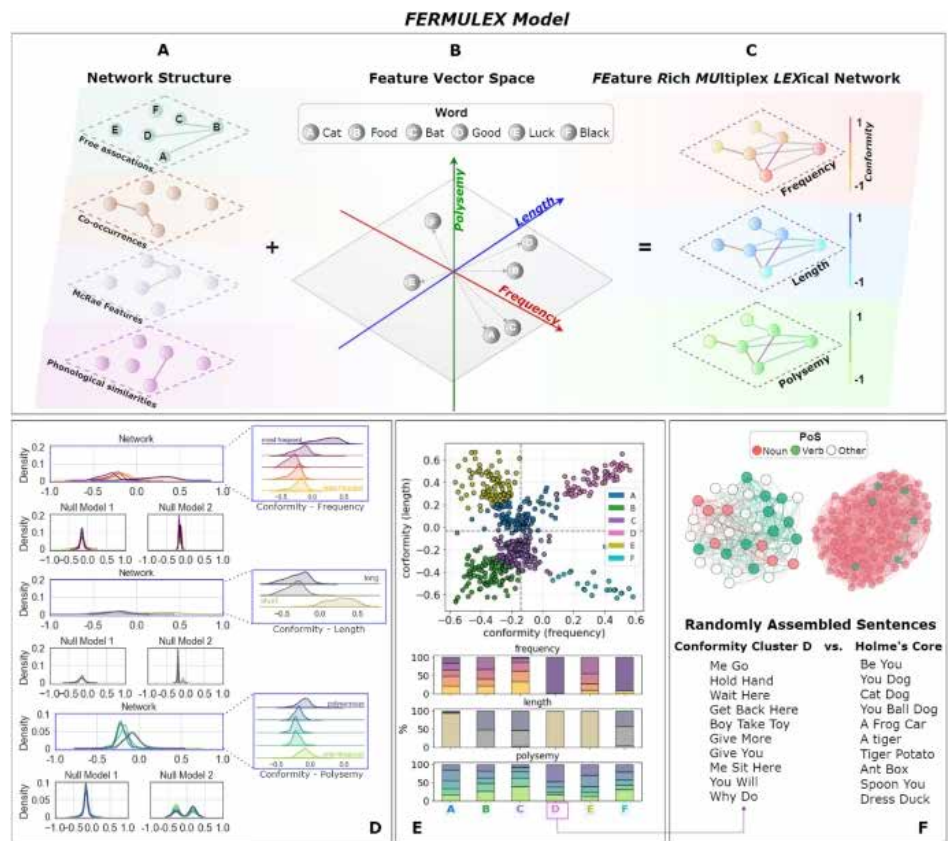
that estimates potential unreported fishing activity hotspots in a marine area. The workflow uses a statistical spatial analysis over vessel speeds and coordinates and a multi-source data integration approach that can work on multiple areas and multiple analysis scales. Specifically, it (i) estimates fishing activity locations and rebuilds data gaps, (ii) estimates the potential unreported fishing hour distribution and the unreported-over-total ratio of fishing hours at a 0.01° spatial resolution, (iii) identifies potential unreported fishing activity hotspots, (iv) extracts the stocks involved in these hotspots (using global-scale repositories of stock and species observation data) and raises an alert about their possible endangered, threatened, and protected (ETP) status. The workflow is also a free-to-use Web Service running on an open science-compliant cloud computing platform with a Web Processing Service (WPS) standard interface, allowing efficient big data processing. As a study case, we focussed on the Adriatic Sea. We reconstructed the monthly reported and potential unreported trawling activity in 2019, using terrestrial AIS data with a 5-min sampling period, containing ~50 million records transmitted by ~1,600 vessels. The results highlight that the unreported fishing activity hotspots especially impacted Italian coasts and some forbidden and protected areas. The potential unreported activity involved 33 stocks, four of which were ETP species in the basin. The extracted information agreed with expert studies, and the estimated trawling patterns agreed with those produced by the Global Fishing Watch.

DOI: 10.3389/fsufs.2023.1152226

Feature-rich multiplex lexical networks reveal mental strategies of early language learning

S. Citraro, M.S. Vitevitch, M. Stella, G. Rossetti
 Scientific Reports, vol. 13. Springer Nature, 2023.

Knowledge in the human mind exhibits a dualistic vector/network nature. Modelling words as vectors is key to natural language processing, whereas networks of word associations can map the nature of semantic memory. We reconcile these paradigms—fragmented across linguistics, psychology and computer science—by introducing FEature-Rich MULTiplex LEXical (FERMULEX) networks. This novel framework merges structural similarities in networks and vector features of words, which can be combined or explored independently. Similarities model heterogenous word associations across semantic/syntactic/phonological aspects of knowledge. Words are enriched with multi-dimensional feature embeddings including frequency, age of acquisition, length and polysemy. These aspects enable unprecedented explorations of cognitive knowledge. Through CHILDES data, we use FERMULEX networks to model normative language acquisition by 1000 toddlers between 18 and 30 months. Similarities and embeddings capture word homophily via conformity, which measures assortative mixing via distance and features. Conformity unearths a language kernel of frequent/polysemous/short nouns and verbs key for basic sentence production, supporting recent evidence of children's syntactic constructs emerging at 30 months. This kernel is invisible to network core-detection and feature-only clustering: It emerges from the dual vector/network nature of words. Our quantitative analysis reveals two key strategies in early word learning. Modelling word acquisition as random walks on FERMULEX



(A–C) Combining multiplex topology (A) and vector spaces (B) results in FERMULEX network (C); (D) kernel density estimates (KDEs) and ridgeline plots highlight conformity distribution for the frequency, length, and polysemy features in toddlers' mental lexicon and the randomised variants; (E) Above—two-dimensional scatter plot of conformity vector space, where each point is colored according to the cluster the point belongs to (K-means algorithm); Below—distribution of word features within each cluster, where a kernel language emerges, i.e. the cluster labeled as D; (F) content characterisation of the kernel compared to a competitor from a k-core decomposition.

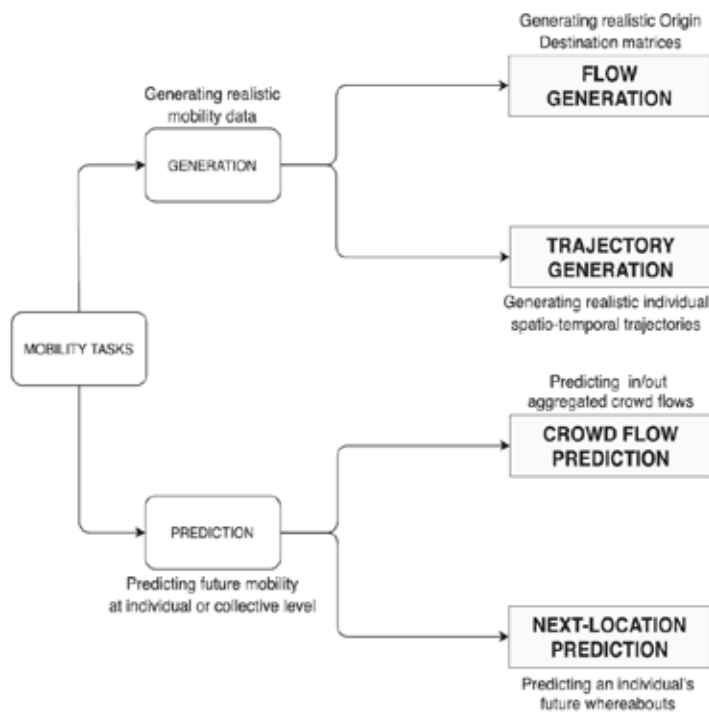
topology, we highlight non-uniform filling of communicative developmental inventories (CDIs). Biased random walkers lead to accurate (75%), precise (55%) and partially well-recalled (34%) predictions of early word learning in CDIs, providing quantita-

tive support to previous empirical findings and developmental theories.

DOI: 10.1038/s41598-022-27029-6

A survey on deep learning for human mobility

M. Luca, G. Barlacchi, B. Lepri, L. Pappalardo
ACM Computing Surveys, vol. 55. ACM, 2023.

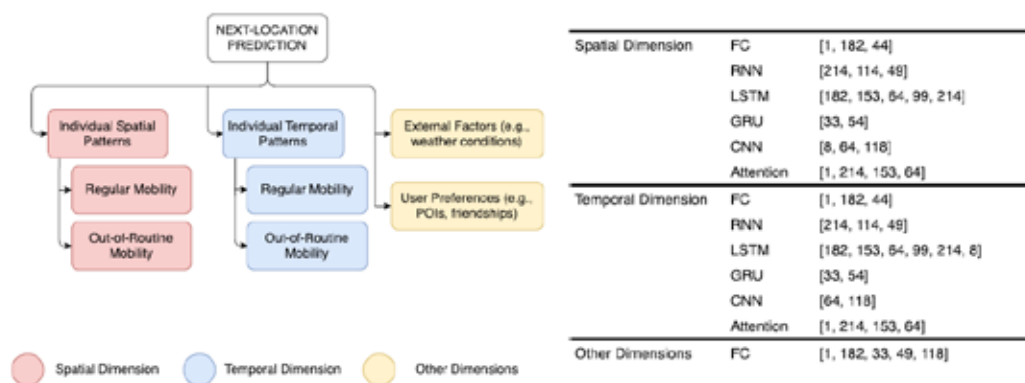


A taxonomy of the mobility tasks we discuss in this survey. Mobility tasks are classified in predictive tasks, aiming at forecasting future mobility at an individual or collective level (Section 3), and generative tasks, aiming at generating realistic trajectories or mobility flows (Section 4). Among the predictive tasks, we cover (i) next-location prediction, the problem of forecasting future whereabouts given the mobility history of individuals (Section 3.1), and (ii) crowd flow prediction, whose goal is to forecast future aggregated flows given historical observations (Section 3.2). However, we have two generative tasks: (i) trajectory generation aims at generating realistic individual trajectories (Section 4.1), and (ii) flow generation aims to generate realistic flows among locations on a geographic region (Section 4.2). We use this taxonomy to map relevant works to the task they solve and shape the survey's structure.

The study of human mobility is crucial due to its impact on several aspects of our society, such as disease spreading, urban planning, well-being, pollution, and more. The proliferation of digital mobility data, such as phone records, GPS traces, and social media posts, combined with the predictive power of artificial intelligence, triggered the application of deep learning to human mobility. Existing surveys focus on single tasks, data sources, mechanistic or traditional machine learning approaches, while a comprehensive description of deep learning solutions is missing. This survey provides a taxonomy of mobility tasks, a discussion on the challenges related to each task and how deep learning may overcome the limitations of traditional models, a description of the most relevant solutions to the mobility tasks described above, and the relevant challenges for the future.

Our survey is a guide to the leading deep learning solutions to next-location prediction, crowd flow prediction, trajectory generation, and flow generation. At the same time, it helps deep learning scientists and practitioners understand the fundamental concepts and the open challenges of the study of human mobility.

DOI: 10.1145/3485125



(a) The aspects next-location predictors should capture regarding the spatial dimension (red), the temporal dimension (blue), and other dimensions (yellow) of human mobility data. (b) DL modules that allow to capture each dimension, with the reference to the selected papers in the literature that implement them.

An optimal algorithm for finding champions in tournament graphs

L. Beretta, F.M. Nardini, R. Trani, R. Venturini

IEEE Transactions on Knowledge and Data Engineering, in press. IEEE, 2023.

A tournament graph is a complete directed graph, which can be used to model a round-robin tournament between n players.

We address the problem of finding a champion of the tournament, also known as the Copeland winner, which is a player that wins the highest number of matches. In detail, we aim to investigate algorithms that find the champion by playing a low number of matches. Solving this problem allows us to speed up several Information Retrieval and Recommender System applications, including question answering, conversational search,

etc. Indeed, these applications often search for the champion inducing a round-robin tournament among the players by employing a machine learning model to estimate who wins each pairwise comparison.

Our contribution allows finding the champion by performing a low number of model inferences. We prove that any deterministic or randomized algorithm finding a champion with constant success probability requires $\Omega(\ln)$ comparisons, where l is the number of matches lost by the champion. We then present an asymptotically optimal determin-


istic algorithm matching this lower bound without knowing l , and we extend our analysis to three variants of the problem. Lastly, we conduct a comprehensive experimental assessment of the proposed algorithms on a question-answering task on public data. Results show that our proposed algorithms speed up the retrieval of the champion up to 13 times with respect to the state-of-the-art algorithm that performs the full tournament.

DOI: 10.1109/TKDE.2023.3267345

Distilled neural networks for efficient learning to rank

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

IEEE Transactions on Knowledge and Data Engineering, vol. 35. IEEE, 2023.



Macro-kernel in the goto algorithm for dense matrix multiplication (DMM).

Recent studies in Learning to Rank have shown the possibility to effectively distill a neural network from an ensemble of regression trees. This result leads neural networks to become a natural competitor of tree-based ensembles on the ranking task. Nevertheless, ensembles of regression trees

outperform neural models both in terms of efficiency and effectiveness, particularly when scoring on CPU. In this paper, we propose an approach for speeding up neural scoring time by applying a combination of Distillation, Pruning and Fast Matrix multiplication. We employ knowledge distillation

to learn shallow neural networks from an ensemble of regression trees. Then, we exploit an efficiency-oriented pruning technique that performs a sparsification of the most computationally-intensive layers of the neural network that is then scored with optimized sparse matrix multiplication. Moreover, by studying both dense and sparse high performance matrix multiplication, we develop a scoring time prediction model which helps in devising neural network architectures that match the desired efficiency requirements. Comprehensive experiments on two public learning-to-rank datasets show that neural networks produced with our novel approach are competitive at any point of the effectiveness-efficiency trade-off when compared with tree-based ensembles, providing up to 4x scoring time speed-up without affecting the ranking quality.

DOI: 10.1109/TKDE.2022.3152585

Raman spectroscopy and topological machine learning for cancer grading

F. Conti, M. D'Acunto, C. Caudai, S. Colantonio, R. Gaeta, D. Moroni, M.A. Pascali
 Scientific Reports, vol. 13. Springer Nature, 2023.

In the last decade, Raman Spectroscopy is establishing itself as a highly promising technique for the classification of tumour tissues as it allows to obtain the biochemical maps of the tissues under investigation, making it possible to observe changes among different tissues in terms of biochemical constituents (proteins, lipid structures, DNA, vitamins, and so on). In this paper, we aim to show that techniques emerging from the cross-fertilization of persistent homology and machine learning can support the classification of Raman spectra extracted from

cancerous tissues for tumour grading. In more detail, topological features of Raman spectra and machine learning classifiers are trained in combination as an automatic classification pipeline in order to select the best-performing pair. The case study is the grading of chondrosarcoma in four classes: cross and leave-one-patient-out validations have been used to assess the classification accuracy of the method. The binary classification achieves a validation accuracy of 81% and a test accuracy of 90%. Moreover, the test dataset has been collected at a different time

and with different equipment. Such results are achieved by a support vector classifier trained with the Betti Curve representation of the topological features extracted from the Raman spectra, and are excellent compared with the existing literature. The added value of such results is that the model for the prediction of the chondrosarcoma grading could easily be implemented in clinical practice, possibly integrated into the acquisition system.

DOI: 10.1038/s41598-023-34457-5

Restoration and content analysis of ancient manuscripts via color space based segmentation

M. Hanif, A. Tonazzini, S.F. Hussain, A. Khalil, U. Habib
 PloS one, vol. 18. Public Library of Science, 2023.



Example of images from DIBCO-2019 [37]: Input degraded image (top row), restored image using the proposed method (bottom row).

Ancient manuscripts are a rich source of history and civilization. Unfortunately, these documents are often affected by different age and storage related degradation which impinge on their readability and information contents. In this paper, we propose a document restoration method that removes the unwanted interfering degradation patterns from color ancient manuscripts. We exploit different color spaces to highlight the spectral differences in various layers of informa-

tion usually present in these documents. At each image pixel, the spectral representations of all color spaces are stacked to form a feature vector. PCA is applied to the whole data cube to eliminate correlation of the color planes and enhance separation among the patterns. The reduced data cube, along with the pixel spatial information, is used to perform a pixel based segmentation, where each cluster represents a class of pixels that share similar color properties in the decorre-

lated color spaces. The interfering, unwanted classes can thus be removed by inpainting their pixels with the background texture. Assuming Gaussian distributions for the various classes, a Gaussian Mixture Model (GMM) is estimated through the Expectation Maximization (EM) algorithm from the data, and then used to find appropriate labels for each pixel. In order to preserve the original appearance of the document and reproduce the background texture, the detected degraded pixels are replaced based on Gaussian conditional simulation, according to the surrounding context. Experiments are shown on manuscripts affected by different kinds of degradations, including manuscripts from the DIBCO 2018 and 2019 publicly available dataset. We observe that the use of a few PCA dominant components accelerates the clustering process and provides a more accurate segmentation.

DOI: 10.1371/journal.pone.0282142

Shallow portion of an active geothermal system revealed by multidisciplinary studies: the case of Le Biancane (Larderello, Italy)

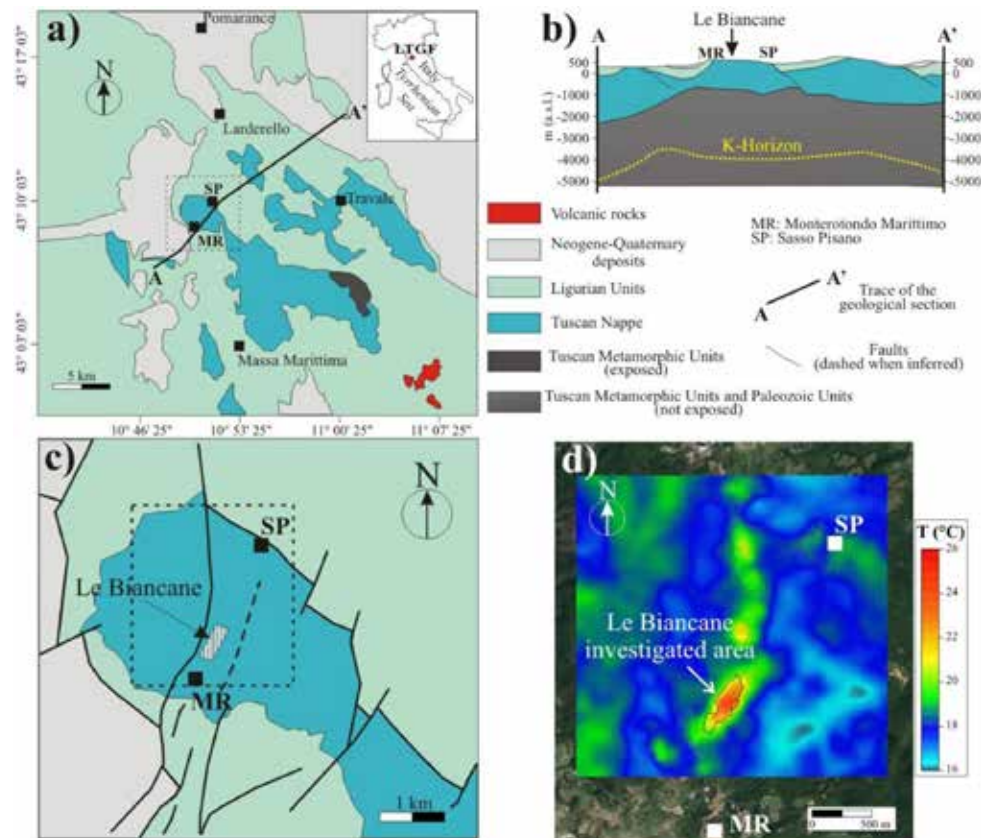
D. Granieri, F. Mazzarini, M. Cerminara, B. Calusi, A. Scozzari, M. Menichini, M. Lelli
Geothermics, vol. 108. Elsevier, 2023.

The natural park of Le Biancane is located in the southern sector of the Larderello-Travale geothermal field (LTGF). It extends over an approximately 100,000 m² area where the impermeable caprock is locally absent and deep fluids may directly reach the surface. Through a multidisciplinary approach including measurements of soil CO₂ flux (total output of 11.5 t day⁻¹), soil temperature

(average 34.4 °C), stable isotope and chemical data on fluids from fumaroles (dominated by a mixture of geothermal gases and air or gases from air-saturated meteoric water), and structural analysis of the formation outcropping, we found that anomalous CO₂ emissions are positively correlated with shallow temperature anomalies. These are in restricted locations adjacent to vents

and fumaroles, where a network of well-connected fractures (preferentially NW-SE and NE-SW orientated and with steep dips) drains efficiently allowing upward migration of the deep fluids and the energy toward the surface.

DOI: 10.1016/j.geothermics.2022.102616

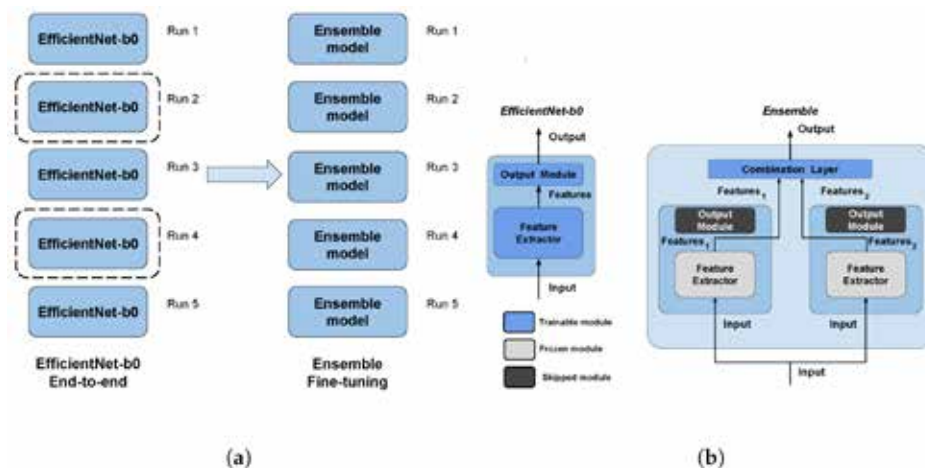


a) Geologic sketch map of the Larderello-Travale Geothermal Field (modified after Liotta and Brogi, 2020). The red dot in the inset is the location of the LTGF; the dotted black box is zoomed in Fig. 1c. b) Geological sketch section (A-A' in Fig. 1a; modified after Arias et al., 2010 and Liotta and Brogi, 2020). c) Geological and tectonic sketch map of Le Biancane area (modified after Liotta and Brogi, 2020); the dotted black box is zoomed in Fig. 1d). d) Map of Land Surface Temperature derived from the scene acquired on 15 June 2021 at 20:59 UTC by the TIRS (Thermal InfraRed Sensor) instrument, onboard the Landsat8 mission, applying the atmospheric corrections described in Barsi et al. (2003). A constant emissivity ($r=0.97$) has been hypothesized for the whole scene, to improve the temperature estimation in the pixels where the land coverage class is mostly bare soil.

Efficient lung ultrasound classification

A. Bruno, G. Ignesti, O. Salvetti, D. Moroni, M. Martinelli


Bioengineering, vol. 10. MDPI, 2023.



The pipeline and the architectures used in this work. The pipeline (a) is made by two main steps: first, end-to-end EfficientNet-b0 training, then ensemble fine-tuning using the best two models (surrounded by a dashed line) of the previous step as weak models. The architectures (b) used are EfficientNet-b0 for end-to-end training, and the ensemble is performed by using a trainable combination layer on the features of the weak models (dark-filled output modules are skipped); moreover, training computational complexity is reduced by freezing the parameters of weak models (light gray filled modules). Both validation steps perform runs with five different seeds (i.e., modules initialization). (a) Validation pipeline. (b) Architectures.

A machine learning method for classifying lung ultrasound is proposed here to provide a point of care tool for supporting a safe, fast, and accurate diagnosis that can also be useful during a pandemic such as SARS-CoV-2. Given the advantages (e.g., safety, speed, portability, cost-effectiveness) provided by the ultrasound technology over other examinations (e.g., X-ray, computer tomography, magnetic resonance imaging), our method was validated on the largest public lung ultrasound dataset. Focusing on both accuracy and efficiency, our solution is based on an efficient adaptive ensembling of two EfficientNet-b0 models reaching 100% of accuracy, which, to our knowledge, outperforms the previous state-of-the-art models by at least 5%. The complexity is restrained by adopting specific design choices: ensembling with an adaptive combination layer, ensembling performed on the deep features, and minimal ensemble using two weak models only. In this way, the number of parameters has the same order of magnitude of a single EfficientNet-b0 and the computational cost (FLOPs) is reduced at least by 20%, doubled by parallelization. Moreover, a visual analysis of the saliency maps on sample images of all the classes of the dataset reveals where an inaccurate weak model focuses its attention versus an accurate one.

DOI: 10.3390/bioengineering10050555



Healthy (first row), pneumonia (second row), and COVID-19 (third row) samples from the dataset and their saliency map. Red (higher) to blue (lower) scale.

NoR-VDPNet++: real-time no-reference image quality metrics


F. Banterle, A. Artusi, A. Moreo, F. Carrara, P. Cignoni
IEEE Access, vol. 11. IEEE, 2023.

Efficiency and efficacy are desirable properties for any evaluation metric having to do with Standard Dynamic Range (SDR) imaging or with High Dynamic Range (HDR) imaging. However, it is a daunting task to satisfy both properties simultaneously. On the one side, existing evaluation metrics like HDR-VDP 2.2 can accurately mimic the Hu-

man Visual System (HVS), but this typically comes at a very high computational cost. On the other side, computationally cheaper alternatives (e.g., PSNR, MSE, etc.) fail to capture many crucial aspects of the HVS. In this work, we present NoR-VDPNet++, a deep learning architecture for converting full-reference accurate metrics into no-reference

metrics thus reducing the computational burden. We show NoR-VDPNet++ can be successfully employed in different application scenarios.

DOI: 10.1109/ACCESS.2023.3263496



The network architecture of NoR-VDPNet (left) and NoR-VDPNet++ (right): Batch Normalization or ReZero are added to each convolution layer.



$\hat{Q} = 86.99$



$\hat{Q} = 86.92$



$\hat{Q} = 56.46$



$\hat{A} = 50.00$



$\hat{A} = 86.92$



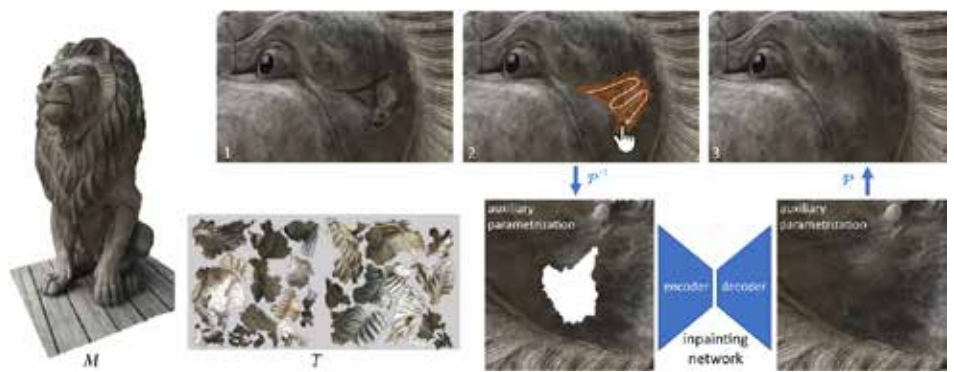
$\hat{A} = 50.00$

An example in which NoR-VDPNet++ is used to choose high-quality images from an image collection. NoR-VDPNet++ predicts a high Q-score (i.e., $Q > 70$) for sharp images, and a low one (i.e., $Q < 60$) for blurred images.

Texture inpainting for photogrammetric models

A. Maggiordomo, P. Cignoni, M. Tarini
Computer Graphics Forum , in press. Wiley, 2023.

We devise a technique designed to remove the texturing artefacts that are typical of 3D models representing real-world objects, acquired by photogrammetric techniques. Our technique leverages the recent advancements in inpainting of natural colour images, adapting them to the specific context. A neural network, modified and trained for our purposes, replaces the texture areas containing the defects, substituting them with new plausible patches of texels, reconstructed from the surrounding surface texture. We train and apply the network model on locally reparametrized texture patches, so to provide an input that simplifies the learning process, because it avoids any texture seams, unused texture areas, background, depth jumps and so on. We automatically extract appropriate training data from real-world datasets. We show two applications




Example of our inpainting operation performed in screen space, over a 3D rendering, introducing artefacts. See also attached video (timestamp: 6 min 38 s).

of the resulting method: one, as a fully automatic tool, addressing all problems that can be detected by analysing the UV-map of the input model; and another, as an interactive semi-automatic tool, presented to the user as a 3D 'fixing' brush that has the effect of

removing artefacts from any zone the users paints on. We demonstrate our method on a variety of real-world inputs and provide a reference usable implementation.

DOI: 10.1111/cgf.14735



A qualitative assessment of the ability of our framework to restore the look-and-feel of a defective texture. From top to bottom: (a) a rendering of a 3D reconstructed model, featuring texture defects caused by missing photographic data; (b) a real photograph of the same object (not used in the texture reconstruction), showing the real aspect of the defected areas; (c) the manually selected region that marks the defected region; (d) the resulting inpainted texture. On the right: two close-ups. Comparing (b) and (d) suggests that in this instance, the overall aspect of the object is successfully recovered.

Dynamic behaviour of the carillon tower in Castel San Pietro Terme, Italy

R.M. Azzara, M. Girardi, C. Padovani, D. Pellegrini

Structural Control and Health Monitoring, in press. Hindawi, 2023.



The Santissimo Crocifisso Sanctuary in San Pietro Terme

The paper describes the experimental investigations conducted on the carillon tower of the Santissimo Crocifisso Sanctuary in Castel San Pietro (Bologna, Italy) and the analysis of data collected by velocimeters and accelerometers installed on the structure. The main goal is to assess the effects of the swinging bells on the dynamic behaviour of the tower. The structure is a rare example of a carillon tower, with fifty-five bells of different sizes, subjected to a careful measurement campaign never carried out before. Six experiments were conducted selectively

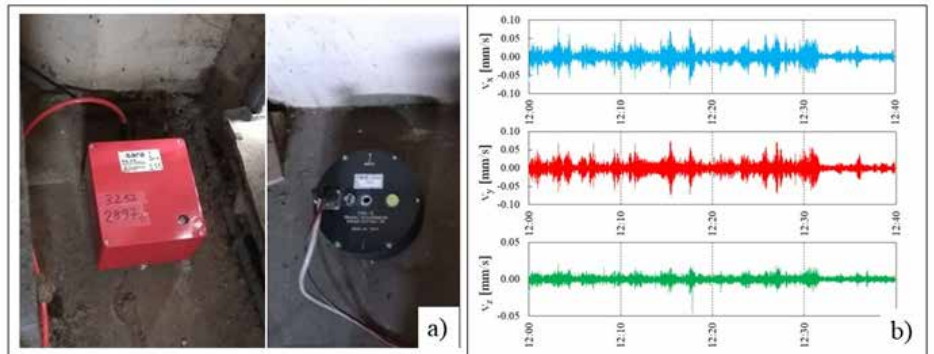
activating the bells, to measure the tower's response induced by different vibration sources and determine the peak velocities recorded by the instruments at different heights. Two ambient vibration tests complemented the six experiments.

The carillon's action induces low velocities on the tower, while experiments involving the bells swinging in the upper chamber produce the highest velocity values in the swinging direction; these values are more

significant than those induced by the carillon alone. The most robust action is induced on the tower when all the bells (carillon plus swinging bells) ring.

The experimental results are complemented by numerical simulations, performed by NOSA-ITACA code (www.nosaitaca.it/software/), of the dynamic behaviour of the tower subjected to the action of a swinging bell.

DOI: 10.1155/2023/1045234



(a) A velocimeter and an accelerometer installed on the bell tower; (b) velocities recorded by using the SS20 2045 seismometer during Experiment 0.

Deep learning for structural health monitoring: an application to heritage structures

F. Carrara, F. Falchi, M. Girardi, N. Messina, C. Padovani, D. Pellegrini

Materials Research Proceedings, vol. 16. Materials Research Forum, 2023.

Thanks to recent advancements in numerical methods, computer power, and monitoring technology, seismic ambient noise provides precious information about the structural behavior of old buildings.

The measurement of the vibrations produced by anthropic and environmental sources and their use for dynamic identification and structural health monitoring

of buildings initiated an emerging, cross-disciplinary field engaging seismologists, engineers, mathematicians, and computer scientists.

In this work, the Authors employ recent deep learning techniques for time-series forecasting to inspect and detect anomalies in the large dataset recorded during a long-term monitoring campaign conducted on the San

Frediano bell tower in Lucca. The problem is framed into an unsupervised anomaly detection task and a Temporal Fusion Transformer is trained to learn the normal dynamics of the structure. The anomalies are then detected by looking at the differences between the predicted and observed frequencies.

DOI: 10.21741/9781644902431-94

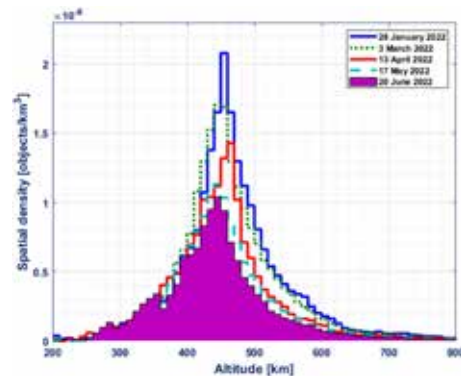
The short-term effects of the Cosmos 1408 fragmentation on neighboring inhabited space stations and large constellations

C. Pardini, L. Anselmo

Acta Astronautica, vol. 210. Elsevier, 2023.

In terms of cataloged debris produced, the anti-satellite test carried out by Russia, in November 2021, at an altitude of about 480 km, leading to the destruction of the old satellite Cosmos 1408, was the second worst to date and represented the third worst fragmentation in orbit. It generated more than 1/4 of the cataloged debris produced over 55 years by all such tests and almost twice as many as were produced by all previous Soviet tests. After placing this event in its historical context, this paper analyzes in detail how the evolution of the Cosmos 1408 debris cloud affected the environment below 600 km in the first seven months, focusing on the two operational space stations and the Starlink large constellation of satellites. During the first six months following the test, the Cosmos 1408 cloud of fragments

nearly doubled the average flux of cataloged objects on the International Space Station

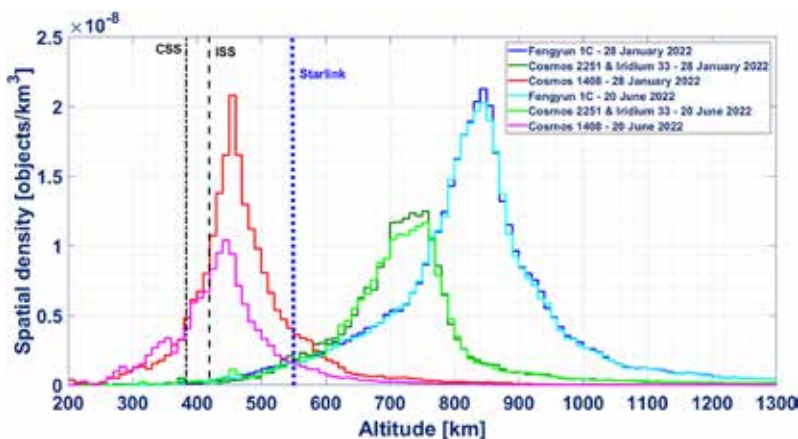
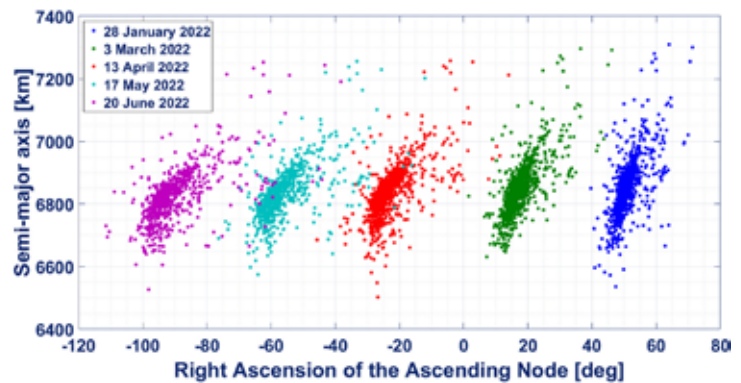


Evolution of the spatial density of the Cosmos 1408 cataloged debris cloud, from January 28 to June 20, 2022. The shrinkage of the cloud due to atmospheric drag, both in the number of objects and in altitude distribution, is evident. The cloud was destined to disappear within a couple of years after the ASAT test.

and increased by about 3/4 that on China's Tiangong. In the same period, the Starlink large constellation saw an average increase in the flux of cataloged objects of about 20%. Some orbital planes, the "counter-rotating" ones with respect to the Cosmos 1408 debris cloud, were more affected than others, and the affected planes gradually changed over time, due to the differential precession of cloud and constellation nodes. However, being the Starlink constellation 70 km higher up, the flux of Cosmos 1408 cataloged debris steadily decreased over the period analyzed, due to the cloud orbital decay, reducing to just over a quarter of its extrapolated initial value after seven months.

DOI: 10.1016/j.actaastro.2023.02.043

Distribution evolution of inclination vs. right ascension of ascending node (mean elements in the True Equator Mean Equinox reference frame) of the Cosmos 1408 cataloged debris.



Spatial density distribution of the cataloged debris clouds generated by the Fengyun 1C and Cosmos 1408 ASAT tests, and by the accidental collision between Cosmos 2251 and Iridium 33.

The VISIONE video retrieval system - runner-up at VBS 2023

Bergen, Norway, January 8-12, 2023



29th International Conference on Multimedia Modeling

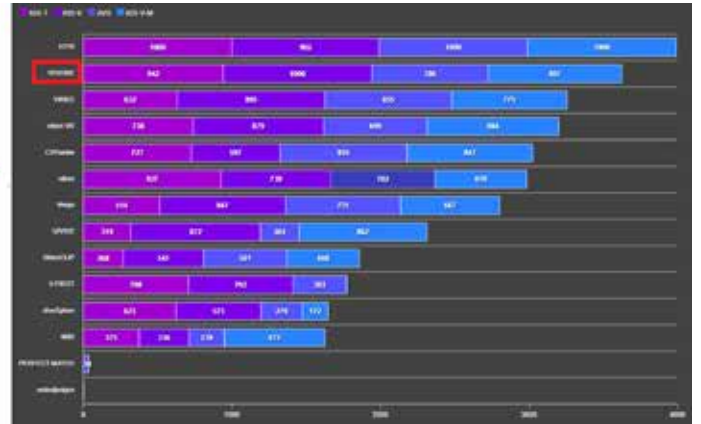
VBS 2023 Runner-Up

presented to

Giuseppe Amato, Paolo Bolettieri, Fabio Carrara, Fabrizio Falchi, Claudio Gennaro, Nicola Messina, Lucia Vadicamo, and Claudio Vairo

for the system

VISIONE at Video Browser Showdown 2023



VISIONE (by G. Amato, P. Bolettieri, F. Carrara, F. Falchi, C. Gennaro, N. Messina, L. Vadicamo, C. Vairo) is a system for fast and effective video search on large-scale datasets, developed at the laboratory for Artificial Intelligence and Media in the Humanities (AIMH).

Recently, VISIONE participated in the 12th Video Browser Showdown (VBS2023), where it ranked highly for most tasks and came second in the overall leaderboard (<https://videobrowsershowdown.org/hall->

[of-fame/](https://videobrowsershowdown.org/hall-)). VBS2023 was held on 9 January, 2023, at the MMM2023 conference in Bergen, Norway. The tasks evaluated during the competition were Known-Item-Search (KIS), textual KIS, and Ad-hoc Video Search (AVS) on about 2300 hours of diverse video content (V3C1+V3C2 dataset) and 12 hours of highly redundant content (Marine dataset). VISIONE integrates several search functionalities that allow a user to search for a target video segment by formulating textual and visual queries, which can also be combined with a temporal search. VISIONE

is based on state-of-the-art deep learning approaches for visual content analysis and exploits highly efficient indexing techniques to ensure scalability. The last version of the system leverages on ALADIN, a new cross-modal, text-to-image retrieval functionality, developed at ISTI-CNR within the AI4Media project.

Contact: Lucia Vadicamo, AIMH Lab
lucia.vadicamo@isti.cnr.it

Thomas Alderighi awarded for the best Ph.D. thesis at University of Pisa

Selected among the best 14 Ph.D. Theses of the University of Pisa 2022

Thomas Alderighi, Ph.D. graduate at the University of Pisa and former member of the Visual Computing Lab of ISTI, is the recipient of the best Ph.D. Thesis Award of the University of Pisa, 5th edition. Every year the University of Pisa organizes a competitive selection to award the best 14 theses defended within the University. In the fifth edition of the prize Alderighi's thesis has

been selected among the best two within the sectors "Mathematics, Computer Science, Physical Sciences, Earth Science".

Thomas has brilliantly defended a thesis on computational methods for improving manufacturing processes leading to top-ranked publications in the Computer Graphics field. During his studies, Dr. Alderighi has been su-

pervised by Paolo Cignoni and Luigi Malomo and is now an R&D Engineer at Ansys Inc.

Contact: Luigi Malomo,
 Paolo Cignoni, VC Lab
luigi.malomo@isti.cnr.it
paolo.cignoni@isti.cnr.it

https://alboufficiale.unipi.it/wp-content/uploads/2023/06/dr88283_atti_settore1.pdf

Gaia Pavoni winner of the VRVis Visual Computing award 2023

Visual Computing Trends Symposium at Tech Gate Vienna, January 26, 2023

Gaia Pavoni received the pan-European VR-Vis Visual Computing Award 2023 (ex aequo with Thomas Höllt from TU Delft) for her research on intelligent digital techniques for monitoring coral reefs. The award was established by the Viennese research centre VRVis Zentrum für Virtual Reality und Visualisierung to acknowledge visual computing solutions that support the achievement of the 17 United Nations Sustainable Development Goals.

With the Visual Computing Lab team, Gaia Pavoni designed an AI-powered open-source annotation tool, TagLab, that greatly simplifies

the work of marine ecologists in the analysis of underwater images. By reducing the time required for ecological interpretation, TagLab enables researchers to process increasingly large volumes of data, ultimately facilitating a greater ability to understand and predict future changes in coral reef ecosystems.

To date, TagLab has been adopted by many important marine science research institutions, such as the Scripps Institution of Oceanography (University of California, San Diego), the Australian Institute of Marine Science (AIMS) and the Mote Marine

Laboratory's Elizabeth Moore International Center for Coral Reef Restoration (Florida, U.S), for the monitoring of reef restoration activities on coral reefs. As an open-source software solution, TagLab aims to mitigate the technological disparities between laboratories. TagLab contributes to the United Nations SDG 13, "Climate Action", SDG 14", "Life Below Water," and SDG 9", "Innovation".

Contact: Gaia Pavoni, VC Lab
gaia.pavoni@isti.cnr.it

The AIMH Lab wins a DH award 2022

DH Awards is an annual international event that recognises projects and initiatives that successfully combine the digital with the humanities. Every year, on the basis of nominations by the community of scholars and practitioners in the Digital Humanities, the DH scientific committee rewards those projects in the different categories that receive the most votes.

The Story Map Building and Visualising Tool (SMBVT), <https://dlnarratives.eu/tool.html>, developed using Semantic Web technologies by the DH group of the Artificial Intelligence for Media and the Humanities (AIMH) Lab, won the DH Awards 2022 in the "Best DH Tool or Suite of Tools" category. SMBVT is a semi-automatic tool that constructs narratives and visualises them in the form of

story maps. Valentina Bartalesi, Emanuele Lenzi and Nicolò Pratelli worked on the the design and development of the tool.

DH Awards results: <http://dhawards.org/dhawards2022/results/>



STORY MAP BUILDING AND VISUALISING TOOL (SMBVT)

Try the [demo version](#) of the tool or ask for credentials to access the [official version](#)

Build a Narrative

about...







or...

Search for a different subject:

Roberto Scopigno awarded the Eurographics gold medal 2023



Hans-Peter Seidel (MPI) and Roberto Scopigno (Cnr-Isti)

9 May 2023, Saarbrücken, Germany. The award recognized the outstanding scientific contribution and leadership of Scopigno in shaping Italian and European computer graphics research. Particular mention was made of his contribution to the Digital Heritage field and the setup of a strong network of collaboration across Europe and the US. Scopigno was also commended for the development of innovative algorithms and techniques in the Cultural Heritage domain and for the acquisition, preservation and visualization of digital copies of physical artifacts inherited from the past.

More details are available at <https://www.eg.org/wp/eurographics-awards-programme/the-eurographics-gold-medal/eurographiss-medal-2023-roberto-scopigno/>


Contact: Roberto Scopigno, Cnr-Isti
roberto.scopigno@isti.cnr.it

Roberto Scopigno (Director of ISTI-CNR) was presented with the Eurographics Gold Medal 2023 (<https://www.eg.org/wp/eurographics-awards-programme/the-eurographics-gold-medal/>), in the opening ceremony of the Eurographics 2023 conference,

AST 2023 1st runner-up paper award

The 4th ACM/IEEE International Conference on Automation of Software Test, May 15-16, 2023 Melbourne, Australia

1st Runner-up Paper Award for the paper "Cross-coverage testing of functionally equivalent programs", by Antonia Bertolino, Guglielmo De Angelis, Felicita Di Giandomenico, Francesca Lonetti.



Cross-coverage increases statement coverage of 26% and branch coverage of 33%

IAASS Vladimir Syromyatnikov - Safety-by-Design award 2023

Award presented to Carmen Pardini and Luciano Anselmo by the International Association for the Advancement of Space Safety (IAASS)



The Space Flight Dynamics Laboratory of ISTI-CNR is the recipient of the **IAASS Vladimir Syromyatnikov - Safety-by-Design Award 2023**. With this award the IAASS honors scientists who have made major technical contributions towards the safety of space systems. The award takes the form

of a bronze statuette representing Athena wielding a silver shield. In Greek mythology, Athena is the goddess of wisdom, science, courage, inspiration, civilization, strength, strategy, justice, skills and shipbuilding. The image on the silver shield represents the APAS docking system designed by Vladimir

Syromyatnikov (1934-2006), which is both a symbol of international cooperation and one of the most successful pieces of space hardware.

Since 1975, the Space Flight Dynamics Laboratory of ISTI-CNR has provided support to national and international space projects, in particular in the field of uncontrolled re-entry. In the last 40 years this Lab has established itself as a global center of excellence in the development of software tools and in the prediction of re-entries. The award recognizes the pioneering work of the Laboratory, and specifically Carmen Pardini and Luciano Anselmo, in the field of re-entry safety. It was presented at the twelfth IAASS conference "Making Space Travel Safer", held in Osaka, Japan, 22-24 May 2023 during the gala dinner at the Sumiyoshi Taisha shrine.

Contact: Carmen Pardini, SFD Lab
carmen.pardini@isti.cnr.it
<https://iaaspace.org/>

Francesco Laccone is the winner of the 2023 Pier Luigi Nervi Prize

At the Italian Workshop on Shell and Spatial Structures (IWSS), Turin 26 - 28 June 2023

Francesco Laccone has won the Pier Luigi Nervi Prize at the IWSS 2023, held in Turin on 26th - 28th June 2023 at the UNESCO World Heritage Site Valentino Castle, Politecnico di Torino.

The Nervi Prize, in memory of the pioneer engineer Pier Luigi Nervi (Italy, 1891-1979),

with the endorsement of the Pier Luigi Nervi Foundation, recognizes talented young researcher working in the field of shell and spatial structures.

The Nervi Prize Commission has selected the winner, examining submissions by 50+ young researchers. To be eligible, the par-

ticipants must be under the age of 35, and be corresponding and presenting author of an IWSS conference paper.

Contact: Francesco Laccone, VC Lab
francesco.laccone@isti.cnr.it
<https://sites.google.com/view/iwss/iwss2023/nervi-prize?authuser=0>

Ph.D. dissertations

Feature-rich networks: when topology meets semantics

Author: Salvatore Citraro, University of Pisa, Department of Computer Science

Supervisor: Giulio Rossetti

Complex networks are a powerful quantitative framework for analyzing complex systems across various domains, including social, biological, and cognitive phenomena, among others. The language of networks allows us to rigorously represent any system as a set of interconnected entities that exhibit emergent behavior and complex-structured patterns, revealing the organizing principles hidden within.

However, real-world data expressed in the form of networks can carry non-structural information that is often ignored by this classic operation of representation. Nonetheless, networks can be enriched with attributes that embed semantic information about entities' properties, like users' politi-

cal values in online social networks, or conceal information encoding different layers of connectivity, such as bus and tram lines in road networks.

Feature-rich networks (see Fig. 1) is an umbrella expression that unifies all these enriched implementations, which can hide behaviors and patterns that classic algorithms are not able to detect. The main goal of this thesis was to provide new algorithms and measures for improving mining on feature-rich networks, such as new methods for network clustering with nodes' attributes, or for (temporal) mixing patterns estimation.

The expressive power of feature-rich networks can be tested in any domain-specific

application where a combination of structural and non-structural mining is essential to unveil the characteristics of a system. In this thesis, the algorithms previously defined were used in many domains to test such effectiveness.

As an example, we report a cognitive network science (CNS) study here. CNS is a field studying human cognition through networked representation of memory and language. To strengthen the importance of the "feature-rich" point of view, we ran processes on a feature-rich representation of cognition (Fig. 2), thanks to which we improved predictions of normative word ordering in children language acquisition.

Bio-inspired approaches for deep learning

Author: Gabriele Lagani (University of Pisa, ISTI-CNR Pisa)

Supervisors: Giuseppe Amato, Fabrizio Falchi, Claudio Gennaro (ISTI-CNR Pisa)

In the past few years, Deep Neural Network (DNN) architectures have achieved outstanding results in several Artificial Intelligence (AI) domains. Even though DNNs draw inspiration from biology, the training methods based on the backpropagation algorithm (backprop) lack neuroscientific plausibility. In my contribution, I considered biologically-inspired solutions for the learn-

ing task. These are interesting because they can help to reproduce features of the human brain, for example, the ability to learn from a little experience. In particular, I explored neuroscientifically grounded Hebbian learning rules, applied to traditional DNNs in combination with backprop, using computer vision as a case study. In Hebbian learning, neurons reinforce their connections when

input and output stimuli are simultaneously excited, and weaken their connections otherwise. Finally, I developed a strategy, named FastHebb, to better exploit GPU acceleration and bring Hebbian approaches to large-scale scenarios.

Computational design & fabrication of tileable patterns: from geometry to mechanical properties

Author: Iason Manolas, University of Pisa, Department of Computer Science

Supervisors: Paolo Cignoni, Luigi Malomo

With the increasing availability of CNC machines and 3D printers, the fabrication of physical artifacts and their visual appearance has become trending research topics in the Computer Graphics community. In recent years, several workflows have been developed to streamline the Digital Fabrication process, overcome material, size, and geometric limitations, and speed up the reproduction and the prototyping phase. In addition to high-resolution reproductions, new approaches which realize objects in an artistic manner instead, have acquired attention. It quickly became ap-

parent that these techniques could also be directed towards the production of objects that look and perform in a desired way, e.g. when subject to a particular external or internal stimulus. In this context, a common theme is the design of ornamental patterns and their use as structural building blocks of complex pattern assemblies bringing into the spotlight the interplay between aesthetics and mechanical properties. In this thesis, we investigate and propose a novel pipeline for designing and efficiently simulating complex pattern tessellations. The thesis presents 3 main contributions.

The first one targets the scarcity of open and efficient simulation tools by proposing a computational tool for predicting the static equilibrium of general bending-active structures which is accompanied by an efficient open-source implementation. Our second contribution is a novel approach for generating a wide range of flat patterns with favorable fabrication-related properties. The third is a computational method for calibrating a reduced mechanical model for each generated pattern enabling the interactive simulation of complex pattern assemblies.

Orchestration strategies for regression testing of evolving software systems

Author: Renan Domingos Merlin Greca, GSSI Gran Sasso Science Institute, L'Aquila

Supervisors: Antonia Bertolino, ISTI-CNR, and Breno Miranda, Federal University of Pernambuco

In today's continuously evolving software, it is desired that changes such as new features and corrections are delivered as quickly as possible. To ensure correct behavior upon release, development teams rely on regression testing suites, which serve to validate previously-correct features and, when well-designed, avoid the propagation of faults to end users. However, regression testing can be costly in large-scale software, and because of this it has undergone extensive research in the last decades.

Notwithstanding, research solutions have hardly found their way into practical usage. To improve applicability of regression testing research, we must identify what are the

main causes of this apparent gap between software engineering academics and practitioners.

We also observe how the largest part of solutions address separately one dimension of the problem at a time. In contrast, we propose a more comprehensive approach that combines techniques with different objectives, including prioritization, selection, reduction and amplification, flaky tests detection, and potentially more. We believe that, by properly orchestrating the differing RT techniques, we can achieve the most from the restricted subset of test cases that can be executed at each new release.

This thesis included the following contributions: a systematic literature review of applicable regression testing research; additional context on the literature provided by the authors of collected primary studies; a test suite orchestration strategy combining robust techniques from the literature; interviews with practitioners at a major technology company that highlight the challenges faced daily by developers and testers; a live repository of papers aggregating relevant literature; a list of challenges that can serve as guidelines for future research.

Posterior probabilities, active learning, and transfer learning in technology-assisted review

Author: Alessio Molinari, Università di Pisa

Supervisors: Fabrizio Sebastiani, Andrea Esuli

Technology-Assisted Review (TAR) refers to the human-in-the-loop machine learning process whose goal is that of maximizing the cost-effectiveness of a review (i.e., the task of labeling items to satisfy an information need). This thesis explores and thoroughly

analyzes: the applicability of the SLD algorithm to TAR scenarios; the usage of active learning combined with the MINECORE framework, effectively improving the framework performance; the portability of machine/deep learning models for the pro-

duction of systematic reviews in empirical medicine. Finally, the thesis proposes a new algorithm, based on SLD, called SALt, which improves the class prevalence estimates on active learning scenarios, with respect to the current state-of-the-art.

Efficiency-effectiveness trade-offs in neural network compression

Author: Cosimo Rulli, Università di Pisa

Supervisors: Franco Maria Nardini, Rossano Venturini

Deep Neural Networks (DNNs) deliver state-of-the-art performance in a huge variety of different fields, from Computer Vision to Natural Language Processing. Their effectiveness comes at the price of huge computational requirements, as modern state-of-the-art DNNs can feature billions of parameters. This poses severe challenges to the usability of DNNs.

In this thesis, we propose three solutions to reduce the computational requirements of DNNs without deteriorating their accuracy in three different problem settings: Learning to Rank (LtR), Image Classification, and multi-term Dense Retrieval.

LtR is the field of machine learning employed to rank candidate documents in a search engine. Recently, neural networks have been applied to LtR, but they struggle to match the tight latency requirements imposed by query processing. In this regard, we employ pruning and cross-modal knowledge distillation to speed up the inference time of neural networks for LtR without

deteriorating their effectiveness. Furthermore, we develop analytic time predictors estimating the execution time of sparse and dense neural networks, thus easing the design of neural models matching the desired time requirements.

For the Image Classification task, we propose Automatic Prune Binarization (APB), a novel compression framework enriching the expressiveness of binary networks with few full-precision weights. APB attains the state-of-the-art accuracy-memory footprint trade-off compared to any other compression method. Moreover, we design two innovative matrix multiplication algorithms for extremely low bits configurations based on highly efficient bitwise and logical CPU instructions. With these routines and sparse-dense matrix multiplication, APB outperforms weight quantization in the efficiency/accuracy trade-off.

In the multi-term Dense Retrieval setting, queries and documents are mapped to a set of vectors by using a language model. The

improved effectiveness compared to single-term representation entails efficiency issues limiting multi-term applications. We propose two different contributions, working with uncompressed and compressed vector representations, respectively. The former exploits query terms and document terms merging to speed up the search phase while jointly reducing the memory footprint. The latter introduces Product Quantization during the document scoring phase and presents a highly efficient filtering step implemented using bit vectors. Our approaches deliver superior efficiency/effectiveness and memory/effectiveness trade-offs compared to existing solutions.

Welcome aboard!



Ilaria Barsanti
Research Staff (Tecnologo)
KDD Lab



Michele Carraglia
Research Staff (Tecnologo)
AIMH Lab



Daniele Fadda
Research Staff (Tecnologo)
KDD Lab



Nicola Messina
Research Staff (Ricercatore)
AIMH Lab



Carlo Metta
Research Staff (Ricercatore)
KDD Lab



Elisa Molinaro
Research Staff (Tecnologo)
InfraScience Lab



Michela Natilli
Research Staff (Tecnologo)
KDD Lab



Roberta Notaris
Admin Staff (Funzionario)



Maria Teresa Paratore
Research Staff (Tecnologo)
InfraScience Lab

Conferences - Co-organized by ISTI



27th ACM International Systems and Software Product Line Conference (SPLC 2023)
28 August - 1 September 2023, Tokyo, Japan

<https://2023.splc.net/>



17th International Workshop on Advanced Infrared Technology and Applications (AITA)
10-13 September 2023, Venice, Italy

<http://aita.isti.cnr.it/>



1st Workshop on AI in Agriculture (AgriAI'23)
17-20, September 2023, Warsaw, Poland

<https://fedcsis.org/sessions/aaia/agriai>



The ACM International Conference on Mobile Human-Computer Interaction
26-29, September 2023, Athens, Greece

<https://mobilehci.acm.org/2023/>



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

Fabrizio Falchi
Silvia Giannini
Carol Peters
Claudia Raviolo
Roberto Scopigno

Layout and Design

Salvatore Rinzivillo, Daniele Fadda, 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