

ERCIM NEWS



Special theme:

AI for Cultural Heritage

Editorial Information

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of about 2,000 printed copies and is also available online, at <https://ercim-news@ercim.eu>.

ERCIM News is published by ERCIM EEIG
BP 93, F-06902 Sophia Antipolis Cedex, France
+33 4 9238 5010, contact@ercim.eu
Director: Dominique Hazaël-Massieux, ISSN 0926-4981

Contributions

Contributions should be submitted to the local editor of your country

Copyright notice

All authors, as identified in each article, retain copyright of their work. ERCIM News is licensed under a Creative Commons Attribution 4.0 International License (CC-BY).

Advertising

For current advertising rates and conditions, see <https://ercim-news.ercim.eu/> or contact peter.kunz@ercim.eu

ERCIM News online edition: <https://ercim-news.ercim.eu/>

Subscription

Subscribe to ERCIM News by sending an email to en-subscriptions@ercim.eu

Editorial Board:

Central editor: Peter Kunz, ERCIM office (peter.kunz@ercim.eu)

Guest Editors: George Pavlidis (Athena RC) and Laura Hollink (CWI)

We gratefully acknowledge their invaluable contributions in curating this special theme on AI in Cultural Heritage. Their expertise and coordination made this issue possible.

Local Editors:

- Ferran Argelaguet, Inria, France (ferran.argelaguet@inria.fr)
- Andras Benczur, SZTAKI, Hungary (benczur@info.ilab.sztaki.hu)
- José Borbinha, Univ. of Technology Lisboa, Portugal (jlb@ist.utl.pt)
- Monica Divitini, NTNU, Norway (divitini@ntnu.no)
- Marie-Claire Forgue, ERCIM/W3C (mcf@w3.org)
- Lida Harami, ICS-FORTH, Greece (lida@ics.forth.gr)
- Athanasios Kalogeras, ISI, Greece (kalogeras@isi.gr)
- Georgia Kapitsaki, Univ. of Cyprus, Cyprus (gkapi@cs.ucy.ac.cy)
- Annette Kik, CWI, The Netherlands (Annette.Kik@cwi.nl)
- Hung Son Nguyen, Univ. of Warsaw, Poland (son@mimuw.edu.pl)
- Alexander Nouak, Fraunhofer-Gesellschaft, Germany (alexander.nouak@iuk.fraunhofer.de)
- Laura Panizo, University of Malaga (laurapanizo@uma.es)
- Erwin Schoitsch, AIT, Austria (erwin.schoitsch@ait.ac.at)
- Thomas Tamsier, LIST, Luxembourg (thomas.tamsier@list.lu)
- Maurice ter Beek, CNR-ISTI, Italy (maurice.terbeek@isti.cnr.it)

Cover photo: View of Guinigi Tower in Lucca, Italy (source: Wikimedia). See also the article "Structural Monitoring of Heritage Buildings Via Deep Learning Algorithms" on page 13.

JOINT ERCIM ACTIONS

- Initiative on Diversity and Inclusion in Computer Science at FORTH-ICS**
Interview with Magdalini Chatzaki, FORTH-ICS
- Report on the AIOTI Workshop: Semantic Interoperability for Digital Twins**
by Dave Raggett, Rigo Wenning (ERCIM/W3C) and Peter Kunz (ERCIM)
- ERCIM "Alain Bensoussan" Fellowship Programme**
- Beyond Compliance 2024 - Research Ethics in the Digital Age**

SPECIAL THEME

Introduction to the Special Theme

- AI for Cultural Heritage**
by Laura Hollink (CWI) and George Pavlidis (Athena RC)

AI for Preservation and Monitoring of Cultural Heritage

- The Unique Authenticity Identifier and AI Tools for Provenance Research of Unregistered Cultural Objects**
by Petros Patias (AUTH), Themistocles Roustanis (KIKLO), and Charalampos Georgiadis (AUTH)

- The ENIGMA Project and Informed Decision-Making on Illicit Trafficking of Cultural Objects**
by Charalampos Georgiadis and Petros Patias (AUTH)

- Structural Monitoring of Heritage Buildings Via Deep Learning Algorithms**
by Maria Girardi (CNR-ISTI), Gianmarco Gurioli (University of Florence), and Nicola Messina (CNR-ISTI)

- AI-Powered Federated Digital Twins for Enhanced Cultural Heritage Preservation and Predictive Analytics**
by Gizealew Alazie Dagnaw and Henry Muccini (University of L'Aquila)

AI for Enhancing Museum and Visitor Experience

- MuNDAR: An AI-Driven, Gamified Augmented Reality Framework for Cultural Heritage Exploration**
by Federico Martusciello and Henry Muccini (University of L'Aquila - DISIM)

- Intelligent Extended Reality Cultural Heritage Guide: Enhancing Immersive Experiences**
by Marina Toumpouri, Václav Milata, and Fotis Liarokapis (CYENS - Centre of Excellence)

20 Culturai: Advancing Mixed Reality Art Experiences with Large Language Models

by Dimitrios Koukopoulos (University of Patras, Greece), Christos A. Fidas (University of Patras, Greece), Marios Belk (Cognitive UX GmbH, Germany)

21 Towards more Inclusive Museums Using AI: Guidance in Sign Language

by Dimitrios Kosmopoulos (University of Patras), Evanthia Papadopoulou (Archaeological Museum of Thessaloniki), Antonis Argyros (FORTH-ICS)

AI for Data Extraction and Cultural Heritage Management

23 TCRMQ: An Application for Generating CIDOC-CRM SPARQL Queries from Text by Using Large Language Models

by Michalis Mountantonakis and Yannis Tzitzikas (FORTH-ICS and University of Crete)

24 AI Multilingual Search Platform for EU Audiovisual Media Archives

by Pilar Orero, Chiara Gunella, and Sarah McDonagh (Universitat Autònoma de Barcelona)

26 Taming AI both as Tool and Artefact for Computer Museums

by Christophe Ponsard (University of Namur), Caroline Djambian, and Emrick Poncet (University of Grenoble)

27 Digital Transformation in Legal History Through Automatic Machine Learning Annotation

by Giovanni Puccetti (CNR-ISTI), Vincenzo Roberto Imperia (University of Palermo), and Andrea Esuli (CNR-ISTI)

29 AI-driven Extraction of Structured Data from a Large Newspaper Archive

by Nikos Kontonasios, Yannis Tzitzikas (FORTH-ICS and University of Crete), and Pavlos Fafalios (FORTH-ICS and Technical University of Crete)

31 Automatic Extraction of Regesta for Medieval Latin Text Summarization

by Giovanni Puccetti (CNR-ISTI), Laura Righi (University of Modena and Reggio Emilia), Ilaria Sabbatini (University of Palermo), and Andrea Esuli (CNR-ISTI)

Ethical, Philosophical, and Sustainability Considerations of AI

32 Who Are You, AI? – Artificial Intelligence as a Potential Metaphysical Humiliation of Humanity

by Michael Fleischhacker (University of Applied Sciences Burgenland), Igor Ivkić (University of Applied Sciences Burgenland and Lancaster University), and Friedrich Roithmayr (University of Applied Sciences Burgenland)

34 Green AI's Role in Cultural Heritage and Sustainability

by Enrico Barbierato and Alice Gatti (Catholic University of the Sacred Heart)

35 Estimating Bias in Book Recommender Systems

by Savvina Daniil and Laura Hollink (CWI)

37 Cultural Contexts in AI Design: Addressing Cultural Biases in AI Models and Data

by Javier Ibanez-Guzman, Imane Taourarti, Arunkumar Ramaswamy (Renault Group) and Ayesha Choudhary (Jawaharlal Nehru University)

38 From Toxic Memes to Toxic Symbols and Back - Towards Automating Toxic Meme Analysis

by Davide Ceolin (CWI), Delfina Sol Martinez Pandiani (University of Amsterdam), and Erik Tjong Kim Sang (NL eScience Center)

RESEARCH AND INNOVATION

40 Development of Advanced Intelligent Robot Platform for Industrial Applications

by János Hollósi, Rudolf Krecht, and Áron Ballagi (Széchenyi István University)

ANNOUNCEMENTS / IN BRIEF

39 Professor Position in AI - XM Chair at the University of Cyprus

42 Dagstuhl Seminars and Perspectives Workshops

42 The Importance of Software

43 Vanessa Evers Appointed New Director of CWI

43 European Project Management

43 ENIGMA Project Announcement

NEXT ISSUE

ERCIM News 142, July 2025

Special theme: Swarm Computing

Call for contributions: <https://ercim-news.ercim.eu/call>

Initiative on Diversity and Inclusion in Computer Science at FORTH-ICS

Interview with Magdalini Chatzaki, FORTH-ICS

Dear Magdalini, What role do you play in your organization?

I am an Msc Computer Scientist working as a telecommunication engineer for the Department of Systems and Networks (DSN) within the Institute of Computer Science (ICS) of the Foundation for Research and Technology Hellas (FORTH). I received my degrees from the Computer Science Department of the University of Crete in mid-1990s and have been working at FORTH-ICS for many years. At that time, female students were among a very small minority in the field of computer science in Greece. As a newcomer I was one of only nine women out of approximately 100 students admitted to the CS department that semester. Although things are much better today, there is still a gap to fill in and progress to balance is extremely slow. Despite the global shortage of computer science specialists and the availability of numerous highly paid vacancies, women and other groups remain far less represented among candidates compared to their male counterparts—an issue worthy of discussion.

Since 2021, I have been a member of the Gender Equality and Anti-Discrimination Committee at FORTH. Beyond my scientific and technical work, I am interested in computer science and its social impact—specifically, the connection between the underrepresentation of women and other marginalised groups in computer science and potential discriminations re-enforced by digital transformation.

FORTH-ICS is not an exception; it mirrors global trends regarding female representation in the field. After years of reflecting on these issues, it is now time to implement interventions within the organisation to raise awareness and initiate research activities to accelerate progress. However, considerable effort and dedicated funding are still required, and currently, the Gender Equality and Anti-Discrimination Committee of FORTH operates on a voluntary basis to promote this cultural shift. It is encouraging that we have the support of the Board of Directors and the organisation's leadership. These are critical factors for ensuring the sustainability of these interventions over the long term. I believe that combining policy initiatives with research activities is essential to achieve tangible results.

Why do you think it is important to promote inclusion and diversity in research institutes and universities?

Academia, universities and research centres are the incubators of knowledge, new ideas, innovation, opportunities, and creators of active citizens and professionals. Academia should not only focus on scientific excellence and advanced technology; it should also aim to cultivate socially sensitive scientists and produce ethically responsible technologies. How important is it to nurture the young scientists in an inclusive and diverse environment?

Highly important I would suggest! If we expect to dismantle social stereotypes, the academic environment is an ideal place—at least one of the primary spaces—to intervene. It is important for young people to develop skills in environments where they feel belong, without experience exclusion or discrimination. Opportunities should be provided inclusively. Moreover, given the ongoing and projected shortage of computer science specialists, institutions should welcome and include diverse social groups.

Can you briefly explain some initiatives that your organisation has started to promote diversity and inclusion? Any initiative that you are particularly proud of?

FORTH-ICS is a research institute of excellence committed to equity, diversity, and inclusion. Since 2021 we have actively participated in the Gender Equality and Anti-Discrimination Committee of FORTH, establishing clear policies to address DEI issues, specifically for FORTH and FORTH-ICS. We developed a Gender Equality Plan, follow-up activities, and a proposed protocol to the Board of Directors for managing discrimination and harassment incidents.

Within the framework of the Gender Equality and Anti-Discrimination Committee at FORTH, we have organised several awareness-raising events. Notably, during FORTH's 40th anniversary, we dedicated a session to highlight our female researchers [L1]. On the International Day of Women and Girls in Science, we promoted women and girls of FORTH, showcasing their research activities and achievements [L2, L3, L4].

Additionally, we actively participate in the National Network of Gender Equality and Anti-Discrimination Committees in Research Centres and Independent Research Institutes [L5], for which we organised a webinar on digital accessibility (in Greek) [L6]. FORTH-ICS introduced the “Mind the Gender Gap” initiative during the Researchers' Night 2024 [L7] featuring a short video [L8] and a banner honouring the institute's female personnel. Lastly, FORTH-ICS organised two workshops: one in July 2023 titled “Considerations on Horizon Europe Gender Sensitive Research @FORTH” [L9] and another on February 2025 titled “Gender Issues in Computer Science” [L10] featuring ACM Distinguished Speaker Prof. Letizia Jaccheri as keynote speaker.

Have you faced any challenge in promoting inclusion and diversity?

One challenge is that inclusion and diversity are often viewed as external to scientific fields—as issues that should concern others, not STEM or computer science specifically. Consequently, finding colleagues willing to invest time and energy is difficult. Another issue is that diversity and inclusion are frequently perceived as concerns exclusively affecting female colleagues. There is also criticism from some colleagues who see these initiatives as promoting a certain “gender ideology”, unrelated to the research community.

Changing attitudes within the research community and highlighting that diversity and inclusion foster a cultural shift benefitting the entire research ecosystem is challenging.



Figure 1: Result of a ChatGPT image prompt.

ChatGPT prompt:

“show me a picture of a successful computer scientist

ChatGPT response:

“Here’s an image of a successful computer scientist in a modern office setting, radiating confidence and professionalism. Let me know if you’d like to see a different style or setting!”

Increased diversity will boost creativity, innovation, democratise the workplace, and open new research directions. The whole concept is not simply to increase female representation but to achieve a more inclusive research environment. Diversity and inclusion are broader than gender issues alone; they encompass disability, accessibility, ethnicity, race, religion, power dynamics, research assessment processes, and practices affecting both workplaces and research outcomes. Given that computer science is behind the construction of digital spaces for humanity, our goal should be designing and building diverse and inclusive digital environments—An ambition we have yet to realise fully.

Another obstacle from the limited representation of women is the resulting lack of role and inspiration models who can inspire young women. Moreover, a lack of diversity among role models is itself problematic. The stereotypical image of the computer scientist as “a very smart guy, slightly weird, geeky or nerdy” remains unappealing, I would say, for many young people.

As an illustrative example, on 11 March 2025 I asked ChatGPT to show me a picture of a successful computer scientist, and Figure 1 shows the result.

Maintaining consistent, promotion of diversity and inclusion poses another challenge, especially without a dedicated budget. Currently, FORTH-ICS lacks allocated funding for these activities. Fortunately, we enjoy ongoing support from our institutional administration and Director, and we hope to secure dedicated funding soon. This funding would be crucial for conducting research and initiating public, educational and community-focused activities.

Is there any “mistake” that is important to avoid?

A common mistake is the belief among some members of the research community that diversity initiatives aim to prioritise women over men. Changing this misunderstanding is challenging. The true goal is to eliminate stereotypes within computer science and establish a diverse and inclusive re-

search ecosystem where career development opportunities are equitable.

Are there lessons learned or best practices you would like to share?

We are not yet fully confident in our understanding of or in having foolproof solutions. However, our experience suggests that initiatives and policies should be informed by quantitative and qualitative data, including feedback from all organisational stakeholders. Policies based on such comprehensive input are likely to gain broader acceptance.

Moreover, sustainable results require persistence and long-term commitment, given that social change typically progresses slowly. Sharing knowledge and experiences at national and international levels is fundamental for achieving lasting outcomes. Finally, embedding inclusion and diversity within research questions and establishing communities and networks to discuss challenges and potential solutions are vital. If we treat diversity as a genuine research problem, our scientific community will undoubtedly develop effective solutions.

Links:

[L1] <https://kwz.me/hFM>

[L2] <https://youtu.be/s6gVgk9r2PY?si=rI8HwkrTEA9-51Bx>

[L3] <https://kwz.me/hFQ>

[L4] <https://kwz.me/hFR>

[L5] <https://sites.google.com/csri.gr/gearnnet-en/the-network>

[L6] https://youtu.be/bpDS_VM2JiA?si=P5f_bVto4P0zpgaB

[L7] <https://www.ics.forth.gr/dsn/mind-the-gender-gap>

[L8] <https://kwz.me/hFU>

[L9] <https://www.ics.forth.gr/dsn/gender-considerations>

[L10] <https://www.ics.forth.gr/dsn/workshop-gender-issues>

The interview was conducted by Monica Divitini of NTNU, chair of the ERCIM Human Capital Task Group.

Please contact:

Magdalini Chatzaki, FORTH-ICS

magda@ics.forth.gr

Recommended Reading

- **Software Engineering and Gender: A Tutorial**
<https://dl.acm.org/doi/10.1145/3663529.3663818>
- **Diversity in Computer Science. Design Artefacts for Equity and Inclusion**
<https://kwz.me/hGT>
- **Challenging systematic prejudices: an investigation into bias against women and girls in large language models**
<https://unesdoc.unesco.org/ark:/48223/pf0000388971>
- **Guidelines for Diversity and Inclusion in Artificial Intelligence**
<https://arxiv.org/pdf/2305.12728.pdf>
- **Gender Parity in the Intelligent Age (World Economic Forum report)**
<https://kwz.me/hGa>

Report on the AIOTI Workshop: Semantic Interoperability for Digital Twins

by Dave Raggett, Rigo Wenning (ERCIM/W3C) and Peter Kunz (ERCIM)

ERCIM hosted the AIOTI (Alliance for AI, IoT and Edge Continuum Innovation) Workshop on Semantic Interoperability for Digital Twins in Sophia Antipolis, France, on 5–6 February 2025. The workshop brought together around 50 experts in the field of semantic interoperability, who enjoyed a packed two-day programme featuring 13 keynotes and 15 long and short talks, with ample space for engaging discussions.



Group photo taken in the sunshine on the first day.

The main goal of the workshop was to advance semantic interoperability to support industrial and commercial applications, specifically leveraging digital twin technology. Semantic interoperability allows computer systems to exchange data with clearly defined and shared meaning, which is critical for seamless integration across multiple domains, supporting initiatives like the EU's Green Deal and circular economy.

Key discussions centred around practical solutions to industrial challenges, the effective reuse of semantic technologies, and the importance of controlled vocabularies, taxonomies, and ontologies. Special emphasis was placed on the role of Digital Twins, virtual replicas of physical entities that support prediction, simulation, and automation capabilities.

Highlights from Day 1 included Martin Bauer (NEC Laboratories Europe) presenting on the AIOTI ontology landscape, emphasizing the reuse and quality assessment of ontologies. Cornelis Bouter (TNO) introduced a forthcoming whitepaper on data-to-ontology mapping, presenting various tools categorised into graphical interfaces, configuration languages, and programming languages, each evaluated for usability, scalability, and expressivity.

Keynote presentations addressed challenges associated with semantic interoperability, notably Vladimir Alexiev (Fraunhofer IWU) discussing the use of Asset Administration Shell (AAS) and ECLASS standards in industrial digital twins, highlighting significant challenges in RDF representation and semantic consistency. Sebastian Käbisch (Siemens) discussed the integration of Web of Things (WoT) standards within the AAS framework, aiming to standardise asset interface descriptions. Darko Anicic (Siemens) stressed the importance of standardised semantics for interoperability in automation systems.

The second day focused on generative AI's role in semantic interoperability, vocabulary development best practices, and the tooling needed for vocabulary management and ontology discovery. Martin Bauer (NEC Laboratories Europe) discussed leveraging generative AI for ontology mapping, highlighting ongoing research into the accuracy and practicality of such applications. Best practices were exemplified through the SAREF suite of ontologies presented by Cornelis Bouter (TNO) and GS1's approach to vocabulary management, emphasizing governance and community involvement by Phil Archer (GS1).

Tooling for vocabulary development received attention, with demonstrations of platforms like OntoPortal presented by Naouel Karam (InfAI) and SEMAPTIC introduced by Carlos Pereira (INESC-TEC), showcasing user-friendly interfaces designed to streamline ontology discovery and application. The event concluded with recommendations emphasizing continued collaboration between industry stakeholders, researchers, and standardisation bodies, underscoring the necessity of user-friendly tools, accurate semantic models, and comprehensive governance for sustainable progress in semantic interoperability.

The workshop called for ongoing dialogue and practical collaboration to address key gaps, especially in data mapping, ontology alignment, and the practical deployment of semantic technologies across diverse industrial domains.

Most of the presentations from the workshop are available on the workshop webpage [L2]. A detailed report, including recommendations, will be published on the ERCIM website [L3]. The workshop chairs would like to thank the program committee and all participants for their contributions to making this event successful. They also express their gratitude to Inria for providing the meeting facilities, and to the EU projects Nephele [L4] and SmartEdge [L5] for their support.

Links:

[L1] <https://aioti.eu/>

[L2] <https://kwz.me/hGi>

[L2] <https://www.ercim.eu/publications/strategic-reports>

[L3] <https://nephele-project.eu/>

[L4] <https://www.smart-edge.eu/>

Please contact:

Dave Raggett, ERCIM/W3C
dsr@w3.org

ERCIM “Alain Bensoussan” Fellowship Programme

The ERCIM Postdoctoral Fellowship Programme is one of ERCIM's principal activities. The programme is open to young researchers worldwide and focuses on a broad range of fields in computer science and mathematics.

The fellowship helps promising scientists improve their knowledge of European research structures and networks, and gain insight into the working conditions of leading European research institutions. Fellowships last for 12 months (with a possible extension) and are spent at one of the ERCIM member institutes.

Where are fellows hosted?

Only ERCIM members may host fellows. When an ERCIM member is a consortium, the hosting institute may be any of its constituent organisations. When an ERCIM member is a funding body, the hosting institute may be any of its affiliates. Fellowships are offered according to the needs of member institutes and the available funding. Fellows are appointed either via a stipend (agreement for a research training programme) or a working contract. The contract type and monthly allowance or salary depend on the hosting institute.

ERCIM encourages researchers from academic institutions and those in industry to apply.

“

After receiving my PhD, I was a bit worried about having focused on a narrow field for an extended period. The ERCIM Alain Bensoussan Fellowship gave me the opportunity to step beyond and explore more topical fields. I was hosted by the prestigious institute, CWI. The experience I gained there has broadened my perspective and deepened my dedication to curiosity. I am now privileged to continue my work as a postdoctoral researcher at CWI, with the chance to further collaborate with those outstanding researchers from diverse backgrounds.



Ezgi DEMIRCAN-TUREYEN
Former ERCIM Fellow



Why apply for an ERCIM Fellowship?

The Fellowship Programme enables outstanding young scientists from around the globe to tackle challenging problems at Europe's leading research centres. It also helps to widen personal ties and deepen mutual understanding among scientists. Through the programme, ERCIM fellows can:

- Work with internationally recognised experts,
- Improve their knowledge of European research structures and networks,
- Familiarise themselves with working conditions in leading European research centres,
- Foster cross-fertilisation and cooperation between research groups.

Equal Opportunities

ERCIM is committed to ensuring equal opportunities and promoting diversity. Applicants for a fellowship within the ERCIM consortium are not discriminated against on the basis of race, colour, religion, gender, national origin, age, marital status or disability.

Conditions

Candidates must:

- Have obtained a PhD in the past eight years (before the application deadline), or be in the final year of doctoral study with an outstanding academic record. Proof of the PhD qualification will be required before the fellowship begins;
- Be fluent in English.

Application deadlines

Deadlines for applications are 31 March and 30 September each year.

Since its inception in 1991, more than 800 fellows have participated in the programme. In 2024, eight scientists began an ERCIM PhD fellowship; over the course of that year, 33 fellows were hosted. The Fellowship Programme is named in honour of Alain Bensoussan, the former president of Inria, one of ERCIM's three founding institutes.

<https://fellowship.ercim.eu>

Beyond Compliance 2024 Research Ethics in the Digital Age

Eight recorded presentations from the 2024 Beyond Compliance workshop, held in October 2024 in Budapest, are now available online. Among the presentations are two particularly brilliant keynote speeches that you should not miss: Julian Nida-Rümelin, philosopher and former German Minister of Culture (“Beyond Compliance: Digital Humanism”) and Milad Doueihi, Professor of Digital Humanities at Paris-Sorbonne University (“Beyond Intelligence: Imaginative Computing”). While Nida-Rümelin focused on tracing the philosophical origins of Digital Humanism and describing its challenges through animism

and mechanistic reductionism, Doueihi offered a historical and literary analysis of what we now refer to as thinking machines.

The ERCIM Ethics Working Group is planning a fourth edition of the Beyond Compliance Workshop series in October 2025. Please contact the ERCIM Ethics Working Group Chair Gabriel David (gabriel.david@inesctec.pt).

For more information, see:

<https://www.ercim.eu/beyond-compliance>

Introduction to the Special Theme

AI for Cultural Heritage

by the guest editors George Pavlidis (Athena RC) and Laura Hollink (CWI)

Artificial Intelligence (AI) has become an integral part of the cultural heritage sector. Libraries, archives, and museums are increasingly employing AI to preserve, interpret, and engage with artifacts of historical and cultural significance. As such, the adoption of AI represents a significant leap forward for the sector. However, this progress brings ethical and practical challenges. Cultural heritage artifacts often reflect diverse and complex histories, and developing AI that can fully capture the breadth and nuance of these narratives is no trivial task. The societal role of cultural heritage institutions as trusted sources of information can sometimes conflict with AI systems that may produce biased results when trained on incomplete or skewed datasets. Moreover, many state-of-the-art AI systems operate as “black boxes,” which creates tension with the sector’s emphasis on data curation and provenance. As a result, there is a growing need for AI systems that are tailored to the specific values and requirements of the cultural heritage field.

This issue of ERCIM News explores “Cultural AI” — a domain that not only applies AI within the cultural heritage sector but also draws inspiration from cultural contexts to inform the design of more inclusive AI systems. The intersection of AI and Cultural Heritage offers a unique opportunity: the cultural heritage sector can help guide the development of AI toward more responsible directions. In doing so, it both benefits from technological innovation and contributes meaningfully to shaping the future of ethical AI.

The articles featured in this issue highlight the wide spectrum of applications in the cultural heritage field that benefit from AI, as well as the diverse forms of AI being employed across the sector. Examples range from the use of deep learning for monitoring the structural health of medieval monuments to LLM-driven dialogues in immersive user interfaces. These contributions, reflecting both the technological breadth and the enormous potential of AI in cultural contexts, can be divided into four connected themes.

AI for Preservation and Monitoring of Cultural Heritage

AI is playing an increasingly central role in the preservation and monitoring of cultural heritage, helping to safeguard historical sites and artifacts from both natural and human-induced threats. Patias et al. introduce the Unique Authenticity Identifier (UAI), a non-destructive AI-based tool for provenance research, designed to track unregistered cultural objects by comparing them to known items and enhancing traceability, particularly in the context of illegal trafficking (page 10). Georgiadis et al. present the ENIGMA project, which integrates AI and machine learning to combat illicit trafficking of cultural goods by providing tools for the identification, provenance research, and traceability of cultural artifacts. Their work emphasizes the importance of data integration and decision-support systems for heritage protection (page 11). Girardi et al. focus on the application of deep learning algorithms for monitoring the structural integrity of heritage buildings. Their approach uses AI to predict potential damage, improving early detection and proactive conservation strategies (page 13). Finally, Dagnaw et al. discuss the integration of Federated Digital Twins (FDT), an AI-powered technology that enables decentralised, multi-institutional collaboration for real-time monitoring and predictive analytics of cultural heritage sites. This innovative approach ensures that data privacy is maintained while fostering collaborative preservation efforts across institutions (page 15).

AI for Enhancing Museum and Visitor Experience

AI is also enhancing how visitors engage with cultural heritage, creating more personalised, immersive, and interactive experiences. Martusciello et al. (page 16) present MuNDAR, a gamified augmented reality (AR) framework that combines AI and AR to enhance visitor engagement in museums and historical sites. By integrating generative AI, MuNDAR provides personalised narratives, challenges, and educational mini-games that adapt to individual visitor profiles. Toumpouri et al. (page 18) explore the use of AI-driven Extended Reality (XR) to create intelligent virtual guides for heritage sites, offering personalised, real-time learning experiences. The Intelligent XR Assistant can adapt content based on visitors' interests and provide them with dynamic, context-aware information. Koukopoulos et al. (page 20) describe CulturAI, a Mixed Reality (MR) application that integrates AI-driven conversational assistants to offer personalised tours of virtual art exhibitions. Visitors can interact with AI guides to learn about artworks, artists, and historical contexts, thus

enhancing their engagement with digital art. Kosmopoulos et al. (page 21) focus on the SignGuide project, which utilises AI to make museums more accessible to Deaf and Hard-of-Hearing (DHH) visitors. The system provides AI-powered sign language guidance, allowing DHH visitors to explore museum exhibits interactively and independently, thus bridging the accessibility gap in cultural heritage settings.

AI for Data Extraction and Cultural Heritage Management

AI is revolutionising the management and analysis of cultural heritage data, enabling institutions to process vast amounts of information more efficiently. Mountantonakis et al. (page 23) describe TCRMQ, an AI tool that uses Large Language Models (LLMs) to generate CIDOC-CRM SPARQL queries from natural language text, allowing for easier querying of cultural heritage data in semantic web formats. Orero et al. (page 24) introduce an AI-based multilingual search platform for audiovisual media archives, which enhances access to European cultural heritage by overcoming language barriers and improving search capabilities across diverse datasets. Ponsard et al. (page 26) discuss the dual role of AI in computer museums, where it is both a curatorial tool for enhancing museum operations and an artifact in itself, reflecting on the history of technology. Puccetti et al. (page 27) describe their work on the automatic extraction of medieval Latin regesta, enabling AI-driven summarisation of historical texts. Their method provides valuable insights for scholars by automating the extraction of concise summaries from lengthy Latin documents. Kontonasios et al. (page 29) present the AI-driven extraction of structured data from historical newspaper archives, using AI tools to digitise and organize historical records, making them more accessible for academic research. Finally, Puccetti et al. (page 31) return with a focus on the REVERINO dataset, which enables AI-driven summarisation of Latin historical texts, demonstrating the potential of AI to support digital humanities research through automated text processing.

Ethical, Philosophical, and Sustainability Considerations of AI

As AI becomes more integrated into cultural heritage practices, it raises important ethical, philosophical, and sustainability concerns. Fleischhacker et al. (page 32) explore the philosophical implications of AI, questioning whether advancements in AI could lead to a “metaphysical humiliation” of humanity, challenging traditional notions of consciousness and

free will. Barbierato et al. (page 34) address the environmental impact of AI, proposing Green AI strategies to reduce the carbon footprint of AI technologies in cultural heritage applications. They advocate for energy-efficient AI practices, such as transfer learning, to mitigate the environmental cost of large-scale AI models. Daniil and Hollink (page 35) highlight the challenges that public media, in particular libraries, face, when implementing a personalised recommender system. They show how popularity bias — a form of bias that is well known to exist in several recommender algorithms — leads to social bias. Taourarti et al. (page 37) examine the cultural biases inherent in AI systems, specifically in the context of autonomous vehicles. They highlight how AI systems trained in one cultural context may perform poorly when deployed in regions with different driving behaviours and infrastructure, stressing the importance of culturally adaptive AI models. Ceolin et al. (page 38) present a framework for analysing toxic memes, exploring how AI can help detect and classify harmful content in digital media. Their work highlights the importance of developing AI systems that can detect hate speech and other toxic content in online platforms, addressing the ethical challenges posed by AI's influence on social media.

Closing Thoughts

AI offers transformative opportunities for the cultural heritage sector, enhancing the preservation, interpretation, and accessibility of cultural assets. However, as these technologies continue to evolve, it is essential to address the ethical, philosophical, and environmental challenges they present. The contributions in this issue demonstrate the immense potential of AI to revolutionise cultural heritage practices, while also underscoring the need for responsible development and deployment of these technologies. By balancing innovation with ethical considerations, the cultural heritage sector can harness the power of AI to safeguard and enrich our shared history for future generations.

Please contact:

Laura Hollink, CWI, The Netherlands
l.hollink@cwi.nl

George Pavlidis, ATHENA RC, Greece
gpavlid@athenarc.gr

The Unique Authenticity Identifier and AI Tools for Provenance Research of Unregistered Cultural Objects

by Petros Patias (AUTH), Themistocles Roustanis (KIKLO),
and Charalampos Georgiadis (AUTH)

The Unique Authenticity Identifier is a novel, non-destructive, multi-parametric, digital indicator that is being developed during the implementation of the EU-funded project ENIGMA [L1]. The UAI will be used to provide a unique provenance metric for unknown/unregistered cultural items by exploiting contextual similarities between the unknown and known items.

The EC Communication on the EU Action Plan against Trafficking in Cultural Goods (2022) [L2] highlighted among others:

- 1) “The identity, authenticity, provenance, and legal status of a cultural good, are rarely instantly visible, and often require specific expertise”;
- 2) “Currently, the available intelligence picture does not adequately reflect the prevalence of cultural goods trafficking”;
- 3) “Moreover, unregistered cultural goods, especially from excavation sites or uncatalogued collections, are hard to detect and track once they have become subject to trafficking.”. In order to tackle these bottlenecks, we need to increase the traceability of cultural goods by leveraging data, increase the capabilities of experts through access to more information, and decrease the response time by automating procedures, while keeping the initial data input at a minimum.

While many cultural goods repositories/databases exist, they are not interoperable, they are not open for sharing, and the information they provide is fragmented. Finally, unknown or fragmented cultural goods lack documentation and are not registered in databases. To overcome these challenges, novel concepts need to be developed. The Unique Authenticity Identifier

(UAI) is focused on identifying undocumented cultural goods by extending the Object ID™ [1] and utilising Artificial Intelligence and Machine Learning tools to provide insights into their provenance, while considering alerts from pinpointed looted sites/monuments. The UAI is one of the core developments of the ENIGMA platform [2].

The UAI uses a 3-tier data classification for feature categorisation. The first tier contains the primary parameters (Subject (general context, e.g., saint, women, men), Type of object (general description, e.g. painting, statue, sculpture, coin), Materials & techniques (broad material classes, e.g. wood, metal, stone), Dimensions, including weight (easily measured in the field), Object photos (captured in the field), Inscriptions & markings, Distinguishing features)). The second tier contains the secondary parameters (Content semantics, Textual content, Historical context (e.g. date, period, age), Spatial/Location associations (e.g. origin, similar locations)). The third tier contains the optimum/miscellaneous parameters (3D characteristics, Image characteristics, Surface/texture characteristics).

The UAI (Figure 1) resembles a DNA strand and, as such, it can be considered the DNA of a cultural object. Its building blocks are the properties / values of the categories, the hierarchical information that originates from standardised vocabularies, the objects image and the related extracted features, and finally the object’s 3D features.

The construction/definition of the UAI is a dynamic process that involves automatic procedures. At the initial stage the Law Enforcement Agency (LEA) officers (Figure 2) begin populating the UAI with basic information, like the object dimensions, its material, its type, and images. All information is recorded, and a ticket is assigned to a mandated expert for further investigation. The ENIGMA platform uses advanced similarity metrics to search in relevant databases, and further populates the UAI, using the information of the documented objects and their similarity with the undocumented one. At the final stage, the expert validates the features that have been added automatically and enhances the information by filling in the remaining blanks based on their expertise. During this step the expert can use and exploit information from ENIGMA's alerts infrastructure. The alerts include information about illegally excavated areas, or looted monuments. The alerts are produced by ENIGMA's earth observation toolkit and crowd-

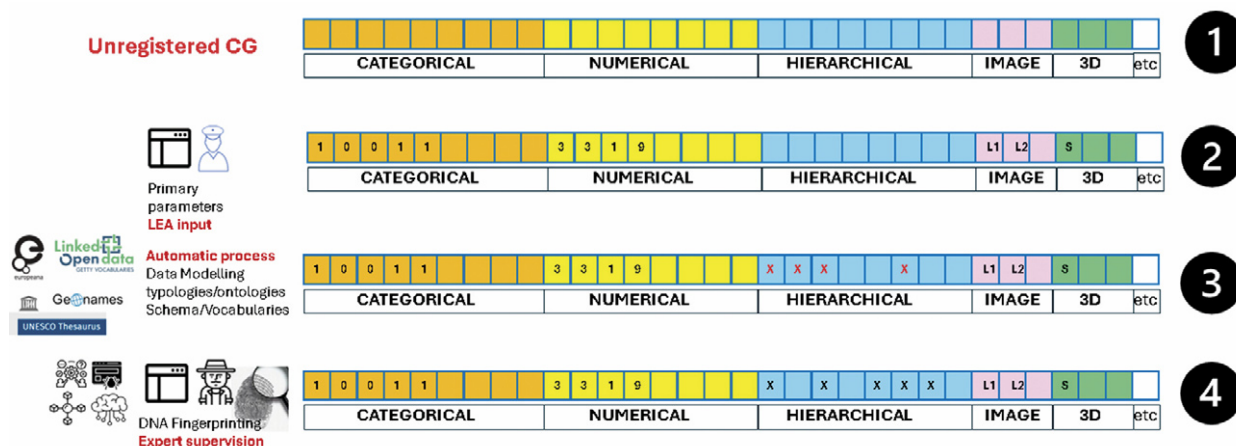


Figure 1: The UAI concept.

Figure 2: CG documentation by a LEA officer example.

sourcing tools. The earth observation toolkit monitors specific areas, selected by mandated authorities, and issues an alert when a potential illegal excavation is detected. The crowd-sourcing tool enables the public to report illegal excavations or looting of monuments. Each step of the procedure is documented and timestamped in order to track the changes made to the UAI and to integrate the dynamic aspect of the data.

The UAI methodology aims primarily to track unknown criminal groups, and to increase the capacity for tracking them by expanding the evidence base. The UAI aims to provide to LEAs and experts with access and the ability to consult the best, worldwide information through AI technologies, and it can be effectively adapted and integrated into their operational workflows and decision-making processes. Finally, the UAI initiates the dialog among stakeholders by highlighting existing gaps, challenging the status quo, rethinking long-term practices, and fostering a culture of continuous learning.

This research was funded by the European Union, under the project ENIGMA “Endorsing, safeguarding, protection, and provenance management of cultural heritage” grant number 101094237. Funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

Links:

[L1] <https://eu-enigma.eu/>

[L2] <https://kwz.me/hFX>

References:

[1] Icom, (1999), Object ID, Available at: <https://icom.museum/en/resources/standards-guidelines/objectid/>.

[2] P. Patias, C. Georgiadis, “Fighting Illicit Trafficking of Cultural Goods—The ENIGMA Project”, *Remote Sensing*, 2023; 15(10):2579. <https://doi.org/10.3390/rs15102579>.

Please contact:

Petros Patias, Charalampos Georgiadis
The Aristotle University of Thessaloniki, Greece
patias@auth.gr, harrisg@auth.gr

The ENIGMA Project and Informed Decision-Making on Illicit Trafficking of Cultural Objects

by Charalampos Georgiadis and Petros Patias (AUTH)

ENIGMA (Endorsing, Safeguarding, Protection, and Provenance Management of Cultural Heritage), is an EU-funded project that aims to achieve excellence in the protection of cultural goods (CGs) and artefacts from man-made threats by contributing to identification, traceability, and provenance research of CGs as well as safeguarding and monitoring endangered heritage sites.

Cultural Heritage (CH) is a testimony to past human activity, and, as such, Cultural Objects exhibit great variety in their nature, size and complexity; from small artefacts and museum items to cultural landscapes, from historic buildings and ancient monuments to city centres and archaeological sites [1].

It has been estimated [2] that 140,000-700,000 antiquities are traded in Europe annually, with a total value of 64M€ – 318M€. It is impossible to distinguish licit from illicit transactions, as there is no way of establishing their authenticity or trading histories.

Generally, the bulk of the trade consists of goods that are small in size and easy to smuggle, such as jewellery and coins. Coins, in particular, take a disproportionate share of seized items. The World Customs Organization (WCO) [L1] reports that the most frequently seized cultural goods are antiquities (such as inscriptions, coins, small seals and the like), followed by archives of sound, film and photographs; household items; and archaeological items. Customs officers identified and recovered 9,399 artefacts (~30% were coins) in 2019, and 22,462 artefacts in 2018 (~50% were coins). In 2016, 69% of seized items were smaller objects, such as antiquities (inscriptions, coins, seals and the like) and historical items (armour, arms). In 2015, 44,235 items were seized and, of those, 98% were coins.

The ENIGMA project [L2, 3] has been running from 1/1/2023-31/12/25, bringing together 12 partners from eight countries to develop a platform that can address, in a holistic way, the illicit trafficking of cultural objects. The platform will integrate, in a unified way, a number of applications and tools that streamline a process for informed decision-making on the illicit trafficking of cultural objects. The platform is intended for the relevant authorities and stakeholders, such as Law Enforcement Agencies and archaeologists or experts working under the mandate of ministries of culture or similar authorities.

The ENIGMA decision support system integrates back-end services with front-end graphical user interfaces (Figure 1) to provide the end user with a holistic tool that assists them in deciding whether or not a cultural good is illicitly trafficked.

The ENIGMA Platform

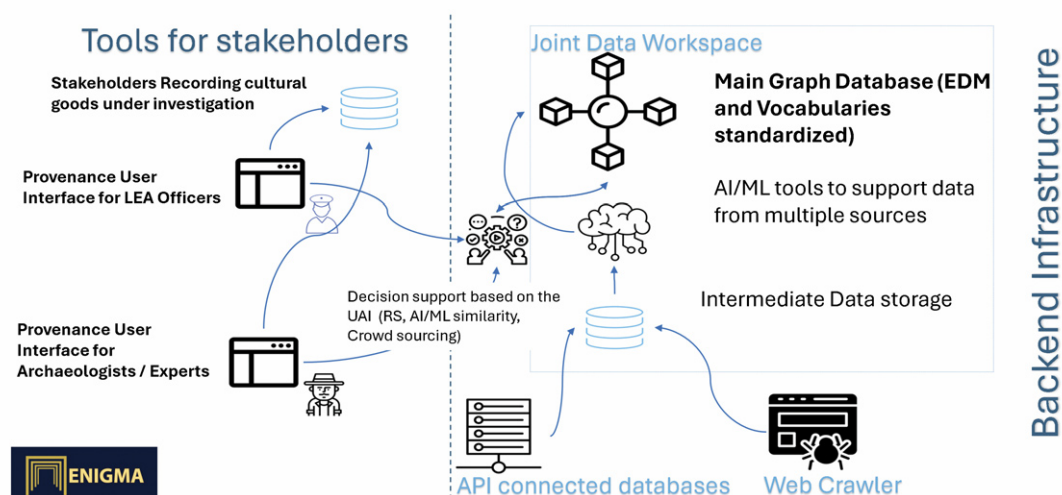


Figure 1: Overview of ENIGMA's platform architecture.

The core of the platform is the joint data workspace, a graph database based on standardised Data Models using standardised vocabularies to offer structured documentation of Cultural Objects.

The main database is populated through API-connected cultural heritage databases, like Europeana. It additionally receives web-crawled data coming from online databases that do not use APIs. Finally, it is populated by the end users when they document a new object. Data coming from unstructured sources are structured and standardised through the use of advanced AI and ML algorithms.

The tools that are integrated in the platform are the following:

- The ENIGMA WebCrawler that periodically searches open-access sites for cultural objects, or sales based on specific sets of keywords and stores the data in a temporary database. The data are then transformed into a standardised structure format and migrated to the ENIGMA Joint Workspace.
- The ENIGMA Earth Observation toolkit, which uses AI and ML algorithms to monitor potential illegal excavation sites using both optical and radar satellite imagery. The areas are

defined by the relevant authorities. The application uses a Geospatial WebGIS infrastructure to select the monitoring areas. Following the area selection, the tool searches for satellite imagery and processes them in near-real-time to detect changes and create alerts on areas where potential illegal excavations have been detected.

- The ENIGMA Crowdsourcing tool that empowers the general public to report potential illegal excavations or looting of buildings through a responsive web application. The users can anonymously report the location, upload images, and add comments to their report. The reports are checked by the responsible authorities and alerts are created if necessary.
- The ENIGMA Scenario Building Engine is a tool aimed for workflow optimization at a higher level that helps the users identify operational or data gaps. This identification can help to optimise and enhance their operational procedures, when encountering illicitly trafficked cultural heritage goods.
- The ENIGMA 3D reconstruction tool uses images or 3D models of fragments to reconstruct the full object. It can provide useful insights to experts about the origin and the provenance of the object.

All these applications contribute to the ENIGMA Provenance Research tool, which is the major front-end application of the platform. The tool is used by Law Enforcement Agency (LEA) officers in the field, for the fast documentation of a cultural object encountered at any travel entry point.

The ENIGMA Provenance Research tool (Figure 2) exploits the data in ENIGMA's joint data workspace to find similarities between unknown/unregistered cultural goods and known / registered cultural goods. The similarity search uses AI/ML techniques to optimise the results leveraging the

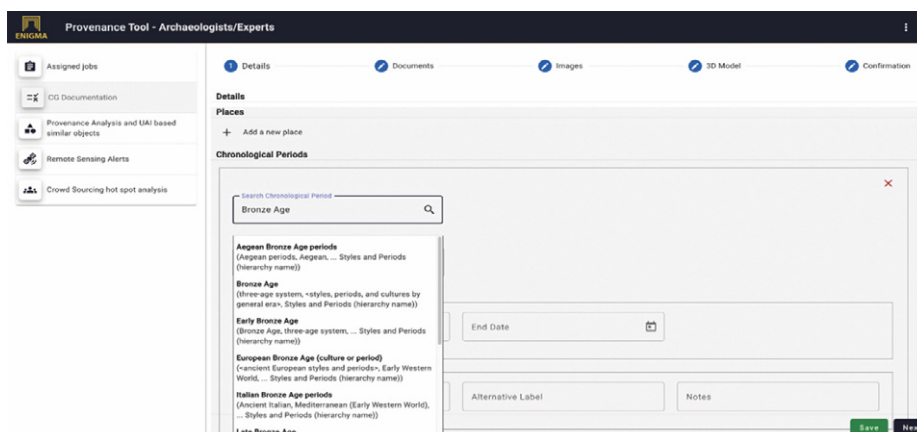


Figure 2: An indicative example of an unknown object's documentation by experts using the provenance research tool.

Unique Authenticity Identifier (UAI), a non-destructive multi-parametric metric. In addition, the alerts created by the Earth Observation toolkit and the Crowdsourcing tool, can provide insights to experts about the origin/provenance of the cultural good.

Currently, the ENIGMA tools are being fine-tuned, optimized, and integrated into the ENIGMA platform and are going to be tested through the implementation of four pilot cases.

This research was funded by the European Union, under the project ENIGMA “Endorsing, safeguarding, protection, and provenance management of cultural heritage” grant number 101094237. Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

Links:

[L1] <https://kwz.me/hFZ>

[L2] <https://eu-enigma.eu>

References:

- [1] P. Patias, “Cultural Heritage Documentation,” in Applications of 3D Measurement from Images, J. Fryer, H. Mitchell, and J. Chandler, Eds. Whittle Publishing Co., 2007, ch. 9, pp. 225–250, ISBN-10: 1-870325-69-9, ISBN-13: 978-1-870325-69-1.
- [2] European Union, “Illicit Trade in Cultural Goods in Europe: Characteristics, Criminal Justice Responses and an Analysis of the Applicability of Technologies in the Combat Against the Trade,” Final Report, ISBN 978-92-79-99359-6, <https://doi.org/10.3390/rs15102579>.
- [3] P. Patias and C. Georgiadis, “Fighting Illicit Trafficking of Cultural Goods—The ENIGMA Project,” Remote Sensing, vol. 15, no. 10, p. 2579, 2023, <https://doi.org/10.3390/rs15102579>.

Please contact:

Charalampos Georgiadis

The Aristotle University of Thessaloniki, Greece
harrisg@auth.gr

Petros Patias

The Aristotle University of Thessaloniki, Greece
patias@auth.gr

Structural Monitoring of Heritage Buildings Via Deep Learning Algorithms

by Maria Girardi (CNR-ISTI), Gianmarco Gurioli (University of Florence), and Nicola Messina (CNR-ISTI)

Monitoring systems constitute a significant, non-invasive tool for verifying the structural health of buildings and infrastructure over time. Deep learning neural networks can be used to analyse data from long-term monitoring systems, such as time series of velocity/acceleration measured at specific points and environmental parameters, and to predict the main features of the buildings' structural behaviour with respect to ambient stresses. Potential anomalies of the structure's vibrational features related to damage or unexpected events, such as earthquakes or exceptional loads, can also be detected. The paper focuses on the application of a Temporal Fusion Transformer (TFT) network to data from the dynamic monitoring of a medieval tower in the historic centre of Lucca (Tuscany, Italy).

The Guinigi Tower (Figure 1) is an iconic medieval monument in the historic centre of Lucca. The dynamic response of this structure to ambient vibrations from the surrounding environment (traffic, wind, earthquakes, crowds and tourists visiting the tower and the nearby monuments) was continuously recorded by a sensor network comprising high sensitivity velocimeters, able to capture small movements at various points of the tower. The experiments lasted about one year, with the tower's velocities recorded at a sampling frequency of 100 Hz. The resulting dataset was used to train a TFT network [1], [L1] and to predict the tower's dynamic behaviour [2]. In particular, the training set was made of the first six natural frequencies of vibration of the tower, extracted from the data every hour of the monitoring period, together with the Root Mean Square (RMS) of the signals and the primary environmental parameters recorded at the tower or in the neighbourhood. The target values are the future frequencies and the RMSs. A change in the frequencies and/or the energetic content of the signal (RMSs) represents an anomaly in the dynamic behaviour of the tower and can be attributed to peculiar events (earthquakes, storms etc.) or exceptional loads. Permanent changes in these quantities can be associated with structural damage. TFT-based monitoring strategy detects when and how much the observed (experimental) values of the targets exceed the confidence interval predicted by the algorithm.

Figure 1 shows the anomaly plots of the first five natural frequencies of the tower (middle) and the RMSs (hourly values) of the signal in the horizontal direction, recorded by one of the seismometers installed over the tower's top. The anomaly plots show the observed values (orange lines), the values predicted by the TFT (blue lines – 50th percentile), and the predicted 1-99th confidence interval (grey areas). The vertical bars indicate when one of the observed targets exceeds the predicted confidence interval, thus indicating an anomaly in our framework. A suitable score is defined to visualise the magnitude of the anomalies via the colour map shown at the bottom.

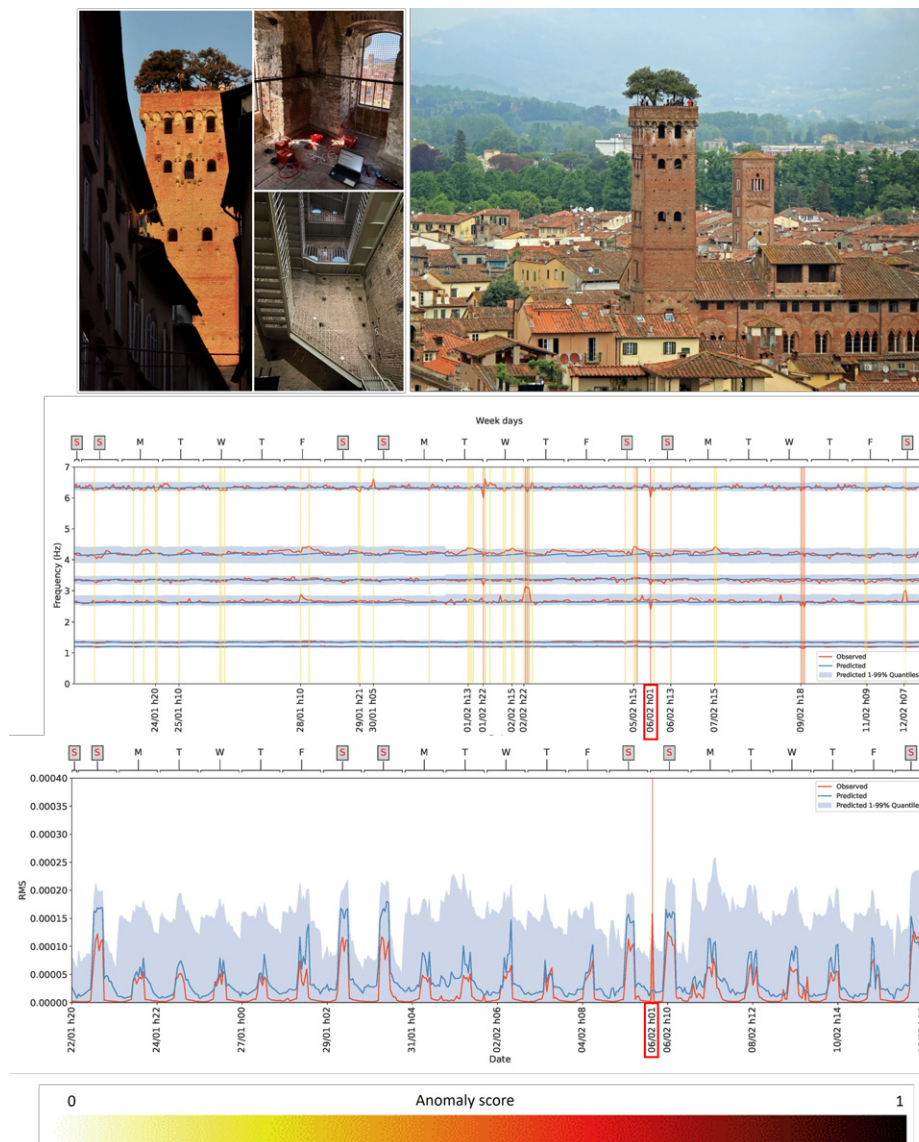


Figure 1: The Guinigi Tower (top). Anomaly plots of the frequencies (middle) and RMSs (bottom). Observed values (orange), predicted values (blue), predicted confidence intervals (grey areas).

The anomaly score, ranging from 0 to 1, quantifies how much the observed target falls outside the confidence interval predicted by the network. Specifically, it reaches the maximum value of 1 when the observed measurement exceeds the estimated interval size of a percentage not smaller than 100%.

The anomaly plot of the frequencies in Figure 1 shows several events, the most significant of which is the Viareggio earthquake (06/02/22, 1:36 UTC, Magnitude = 3.7 [L2]), which happened about 20 km from Lucca. As illustrated by the figure, the anomaly induced by the earthquake is detected by the algorithm in every one of the tower's frequencies and by examining the RMSs. After the event, the values of the frequencies and RMSs came back to normality, thus indicating that the earthquake did not permanently damage the tower's structure. The anomaly plot of the RMSs is also meaningful: the daily increase in the signal's energy, which is well predicted by the algorithm, is related to the presence of the visitors on the tower during the day (the number of tickets sold per hour is included for training the algorithm). The algorithm output correlates well with the increase in weekend visitors, marked in the figure by the Saturday-Sunday capital letters.

The authors also explored the potential of TFT to forecast the dynamic behaviour of heritage buildings using the datasets recorded at the San Frediano Bell Tower in Lucca [3] and the

Matilde Donjon in the Medicean Port of Livorno. Leveraging the network's internal features to automatically classify different anomaly categories and damage levels is a worthwhile future perspective for fully automated applications in structural monitoring.

Links:

- [L1] <https://pytorch-forecasting.readthedocs.io/>
[L2] <https://itaca.mi.ingv.it>

References:

- [1] B. Lim, et al., "Temporal fusion transformers for interpretable multi-horizon time series forecasting", in *Int. J. of Forecasting*, 37(4), pp.1748-1764, 2021.
- [2] F. Falchi, et al., "Deep learning and structural health monitoring: Temporal Fusion Transformers for anomaly detection in masonry towers", in *Mechanical Systems and Signal Processing*, 215, 111382, 2024.
- [3] F. Carrara, et al., "Deep learning for structural health monitoring: An application to heritage structures", in *Materials Research Proc.*, 26, pp. 581–586, 2023.

Please contact:

Maria Girardi, CNR-ISTI, Italy
maria.girardi@isti.cnr.it

AI-Powered Federated Digital Twins for Enhanced Cultural Heritage Preservation and Predictive Analytics

by Gizealew Alazie Dagnaw and Henry Muccini (University of L'Aquila)

Federated Digital Twin (FDT) technology, enhanced by Artificial Intelligence (AI), is revolutionising cultural heritage preservation. By enabling decentralised, collaborative monitoring, predictive analytics, and immersive interactions, FDTs ensure the longevity, accessibility, and sustainability of historical sites. This innovative approach optimises conservation strategies through multi-institutional cooperation, paving the way for future advancements in heritage management.

Cultural heritage sites face growing threats from climate change, pollution, and human activities. Traditional preservation methods often fall short in providing real-time monitoring and predictive capabilities, leading to potentially irreversible damage [1]. Federated Digital Twin (FDT) technology, powered by Artificial Intelligence (AI), offers a transformative solution by creating decentralised virtual models of cultural heritage sites. These models enable real-time insights, predictive conservation, and enhanced public engagement [2]. FDT systems facilitate collaboration among multiple institutions, al-

lowing for data aggregation without compromising sensitive information. Through federated learning and AI-powered models, these virtual twins provide critical insights into site evolution under various conditions, enabling proactive interventions and risk reduction. Furthermore, the integration of real-time environmental data and structural analysis optimises conservation efforts and enhances decision-making [3]. Extended Reality (XR) technologies are also being combined with Digital Twins to create immersive and interactive heritage experiences [1].

This project is conducted under the National Recovery and Resilience Plan (PNRR) at the University of L'Aquila, Italy, fostering collaboration among multiple stakeholders including museum administrations, software engineering research groups, and AI research groups. Running from November 2024 to December 2025, this interdisciplinary effort aims to enhance cultural heritage preservation through innovative technological solutions. The objective is to investigate AI-powered Federated Digital Twins for cultural heritage preservation, with a focus on decentralized data collaboration, automated monitoring, predictive conservation, and interactive heritage experiences.

Federated Digital Twin Approach

The core innovation of this project lies in the application of Federated Digital Twin (FDT) technology. Through FDT, multiple heritage institutions can collaborate on preservation efforts without sharing sensitive data [2]. This approach aggregates localised insights from various sites, enabling the development of robust and accurate predictive models for conservation. FDTs are particularly suited for monitoring cultural heritage sites across different institutions while maintaining data privacy. For instance, federated learning algorithms allow cultural heritage sites in different regions

to share insights on deterioration patterns, structural vulnerabilities, and environmental risks, facilitating the development of adaptive preservation strategies. The decentralised nature of FDTs also supports efficient, large-scale data analysis while minimizing computational overhead.

Contributions from the Federated Digital Twin Perspective

Our research aims to optimise the application of FDT technology in real-world cultural heritage settings. The primary contribution is the integration of FDTs for predictive analysis and conservation decision-making across multiple museums in L'Aquila, Italy. The project specifically simulates the MuNDA (Museo Nazionale d'Abruzzo), Amiternum, il Castello (Spanish Fort), and MAXXI L'Aquila museums. By employing federated learning algorithms, we enable these institutions to optimise resource management, mitigate risks, and enhance

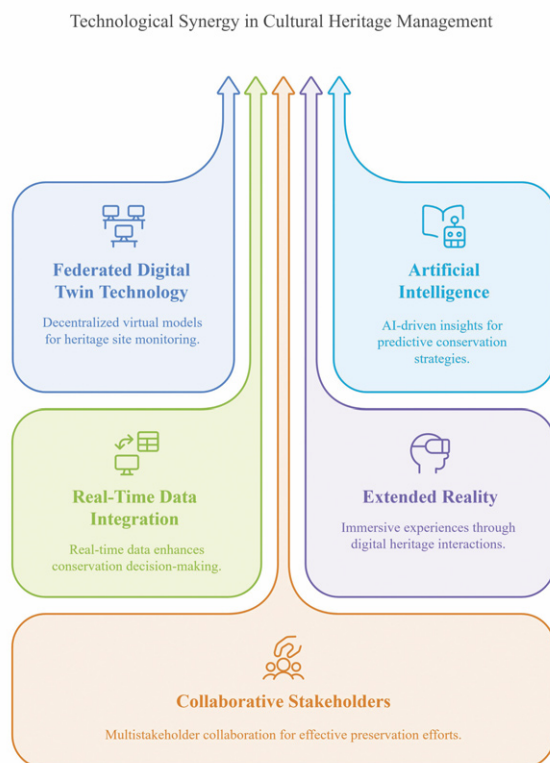


Figure 1: AI-powered Federated Digital Twin framework.

the visitor experience through coordinated strategies based on real-time data and predictive simulations. The Federated Digital Twin approach facilitates continuous improvements as data is aggregated from various sites, resulting in richer and more accurate models for future predictions and intervention strategies.

Conclusion

Federated Digital Twins are poised to revolutionize cultural heritage preservation by enabling decentralised, collaborative decision-making, predictive conservation, and interactive digital reconstructions. These intelligent systems provide a dynamic, data-driven approach to heritage management, enhancing the sustainability and accessibility of historical sites while minimizing the risks associated with traditional preservation methods. By integrating AI-powered models and federated learning, FDTs offer a promising future for the proactive and efficient conservation of cultural heritage.

References:

- [1] F. Bruno, et al., "From 3D reconstruction to virtual reality: A complete methodology for digital archaeological exhibition," *J Cult Herit*, vol. 11, no. 1, pp. 42–49, 2010, doi: 10.1016/j.culher.2009.02.006.
- [2] L. Bonorden et al., "Decision-Making About Federated Digital Twins - How to Distribute Information Storage and Computing," in *Lecture Notes in Informatics (LNI), Proceedings - Series of the Gesellschaft für Informatik (GI), Gesellschaft für Informatik (GI), 2022*, pp. 49–64. doi: 10.18420/modellierung2022-006.
- [3] N. Moretti, et al., "Federated data modelling for built environment Digital Twins."

Please contact:

Henry Muccini, University of L'Aquila, Italy
henry.muccini@univaq.it

Gizealew Alazie Dagnaw, University of L'Aquila, Italy
gizealewalazie.dagnaw@graduate.univaq.it

MuNDAR: An AI-Driven, Gamified Augmented Reality Framework for Cultural Heritage Exploration

by Federico Martusciello and Henry Muccini (University of L'Aquila - DISIM)

The rapid advancement of digital technologies is revolutionizing how we engage with Cultural Heritage (CH). MuNDAR is an ambitious project focused on the integration of Gamification (GAM), Generative Artificial Intelligence (GAI) and Augmented Reality (AR) to create immersive and personalised user journeys in mobile applications in the context of CH. By integrating these technologies, we aim to enhance user engagement, accessibility, and storytelling, offering a dynamic and interactive way to explore history and culture.

The Role of ICT in Cultural Heritage

In recent years, Information and Communication Technologies (ICT) have significantly transformed the CH sector, offering unprecedented opportunities to enhance visitor engagement and experiential value. Institutions worldwide are increasingly utilising digital tools to digitise and enrich their collections while seeking innovative methods to attract wider audiences.

The Department of Information Engineering, Computer Science, and Mathematics at the University of L'Aquila, in collaboration with the Department of Human Sciences and the Museo Nazionale d'Abruzzo (MUNDA), has developed an advanced research prototype integrating Gamification (GAM), Generative Artificial Intelligence (GAI), and Augmented Reality (AR) to redefine visitor experiences at key cultural sites in L'Aquila. Launched in 2023 and scheduled for completion in 2025, this initiative introduces a novel software architecture to facilitate interactive and immersive engagement with cultural heritage. Initial testing has demonstrated promising outcomes, emphasizing the potential of this approach.

The research team, with expertise spanning software architecture, digital humanities, and museology, has designed a reference software framework that seamlessly integrates GAM, GAI, and AR. This architecture ensures smooth interoperability between these technologies, providing visitors with an enriched and cohesive experience.

Architectural Framework and Technological Integration

The implemented system shown in Fig 1 is based on a multi-layered software architecture integrating the following core components:

- Gamification Engine: Enhances engagement by incorporating points, rewards, and challenges tailored to visitor behaviour and interests.
- AR: Provides real-time visualization of historical reconstructions and interactive elements through mobile applications.

- GAI: Generates context-aware, dynamic content, adjusting narratives, pathways, and explanations to suit different user profiles, ensuring accessibility and inclusivity.
- Integration Framework: A middleware layer that facilitates seamless interaction between AI-driven content, AR experiences, and gamified exploration paths.

This architecture allows for a scalable, modular approach, enabling future enhancements and broader applicability across various cultural heritage contexts.

Gamification as a Tool for Cultural Engagement

Cultural heritage sites often face challenges in attracting and maintaining audience engagement. Gamification, defined as the application of game-design elements in non-gaming contexts, has been shown to be highly effective in fostering learning and active participation. By incorporating interactive mechanics such as rewards, quests, and challenges, cultural institutions can create more compelling experiences [1].

This project enhances museum visits through gamified storytelling and exploration. Focusing on three major locations—the MUNDA museum, the Amiternum archaeological site, and the 16th Century Castle of L'Aquila—the research seeks to highlight MUNDA's extensive collections while increasing the visibility of lesser-visited archaeological sites. A digital treasure hunt has been implemented, where users progress by discovering points of interest (POI) and engaging in educational mini-games and challenges linked to them.

Merging the Physical and Digital: The Phygital Approach

Traditional museum experiences often suffer from the decontextualisation of artifacts, as objects are displayed far from their original sites, detaching them from their historical narra-

tives. This project addresses this issue through a “phygital” approach integrating physical and digital dimensions to contextualise artifacts within their original settings.

AR plays a crucial role in bridging this gap by enabling visitors to visualise historical objects in the locations where they were originally discovered [2]. When standing at an archaeological site, users can access AR-driven reconstructions, offering a real-time visual representation of how artifacts were once positioned. This approach not only enhances engagement but also deepens visitors' understanding of historical contexts by merging tangible and intangible heritage elements through storytelling and interactive digital content.

AI-Driven Personalization for Enhanced Cultural Exploration

Gamification mechanics alone are not sufficient to create highly personalised cultural journeys. GAI is leveraged within this project to enable adaptive content generation, decision-making support, and interactive user experiences [3]. AI-powered systems analyse visitor interactions, preferences, and prior knowledge to deliver tailored recommendations and dynamic storytelling elements.

Beyond static reconstructions, GAI creates dynamic narratives that evolve based on user input. Visitors receive customised pathways through cultural sites, with AI adapting content and pathways in real time. This ensures that each visitor's experience is uniquely relevant, enhancing engagement and learning outcomes.

The system also gathers continuous feedback from users, refining its adaptive mechanisms to align with their expectations and needs. By incorporating AI-driven personalization, the

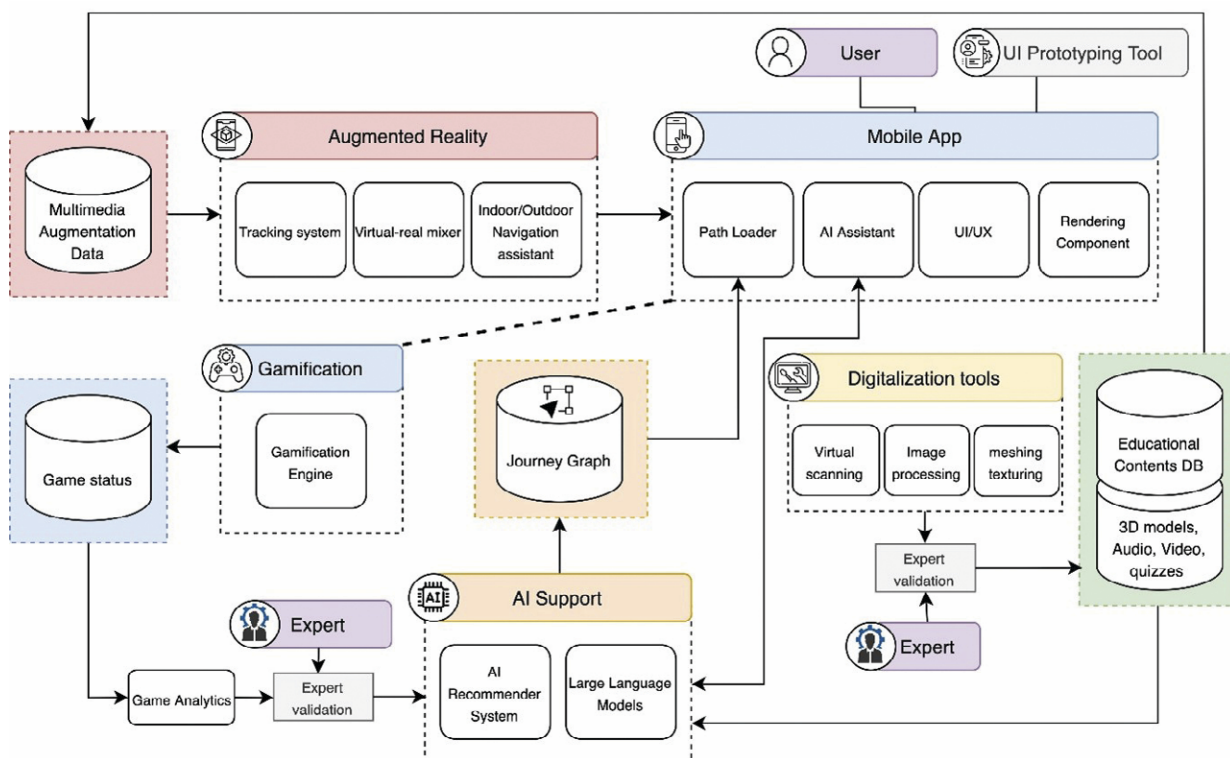


Figure 1: The Implemented architecture. The map is categorised into six distinct colours: digitalization tools by yellow, educational contents repository is represented by green, gamified elements by blue, actors by purple, AI by orange, augmented reality by red.

project ensures a high degree of relevance and accessibility, making cultural exploration more engaging for diverse audiences.

As shown in Figure 1, the GAI module utilises a recommender system, journey path creator, and Large Language Models (LLMs) to generate adaptive content and personalised visitor pathways. The selection of materials by experts is essential for maintaining content quality, as educational specialists curate and validate 3D models, audio, video, and quizzes stored in the educational content database. Additionally, fine-tuning of LLMs ensures that AI-generated content aligns with expert-approved historical and educational accuracy. This guarantees that AI-driven narratives and interactions remain reliable, contextually appropriate, and enriching for visitors.

Future Directions: Expanding the Digital Cultural Experience

The next phase of the project will focus on developing a fully operational system capable of generating complex, real-time personalised visitor journeys without the supervision of a path validation expert. By refining user profiling methodologies and interactive behaviour analysis, the system will dynamically tailor content and routes to match individual visitor interests.

Future efforts will emphasise:

- Enhancing AI-driven content adaptation for improved narrative coherence and educational impact.
- Refining real-time analytics for better responsiveness to visitor interactions.
- Expanding collaborations with cultural institutions to scale the project's impact.

By utilising cutting-edge digital tools, this initiative aims to set a new standard in cultural heritage engagement, transforming traditional museum experiences into dynamic, interactive explorations of history and culture.

References:

- [1] C. Özgül and F. Erbay, "Gamification practices in museums," *Journal of Tourismology*, vol. 7, no. 2, pp. 265–276, 2021. <https://doi.org/10.26650/jot.2021.7.2.1017009>
- [2] Y. Zhou, J. Chen, and M. Wang, "A meta-analytic review on incorporating virtual and augmented reality in museum learning," *Educational Research Review*, vol. 36, p. 100454, Jun. 2022. <https://doi.org/10.1016/j.edurev.2022.100454>
- [3] M. Bugeja and E. M. Grech, "Using technology and gamification as a means of enhancing users' experience at cultural heritage sites," in *Proceedings of the International Conference on Digital Cultural Heritage*, D. Seychell and A. Dingli, Eds., 2020. https://doi.org/10.1007/978-3-030-36107-5_4

Please contact:

Federico Martusciello, University of L'Aquila (DISIM), Italy
federico.martusciello@graduate.univaq.it

Henry Muccini, University of L'Aquila (DISIM), Italy
henry.muccini@univaq.it

Intelligent Extended Reality Cultural Heritage Guide: Enhancing Immersive Experiences

by Marina Toumpouri, Václav Milata, and Fotis Liarokapis
(CYENS - Centre of Excellence)

AI-powered chatbots and virtual guides are transforming cultural heritage, making it more interactive, immersive, and accessible. Leading this transformation, the EU-funded XR4ED project is driving innovation in XR-based education and training. A key outcome of the project is the Intelligent XR Assistant, an AI-driven system enhancing user interaction, adaptive learning, and real-time engagement in XR environments. As part of this effort, an AI-powered guide for Nicosia brings its historic centre to life with AI avatars. By leveraging advanced XR and AI technologies, the assistant transforms the way cultural heritage is experienced.

The emergence of artificial intelligence (AI) has significantly transformed the way cultural heritage is perceived, experienced, and consumed, as it has the capacity to enhance visitor experiences through personalised, immersive, and interactive technologies [1]. Notably, AI-powered chatbots and virtual guides can provide real-time, multilingual information about historical sites, museums, and landmarks. AI-driven augmented and virtual reality (AR/VR) facilitate immersive experiences of historical environments, including lost architectural heritage, thus providing deeper insights into past civilisations and cultures. Complementing this, AI-driven recommendation systems can tailor itineraries based on visitors' age, preferences, educational level, and other factors, thereby enriching their interaction and engagement with cultural heritage. Furthermore, machine-learning algorithms can contribute to the analysis and digital preservation of artefacts, facilitating the creation of virtual exhibitions and high-quality 3D models that can be shared globally. By leveraging AI, these advancements make cultural heritage tourism more engaging and educational, while effectively bridging the past with modern technology.

XR4ED [L1] is a project funded by the European Union, designed to foster collaboration between the EdTech and XR (Extended Reality) communities, with the objective of establishing a centralised European platform for XR-based learning and training. This initiative seeks to enhance digital education by providing access to XR content, tools, and solutions, thereby supporting start-ups, SMEs, and educational institutions in developing market-ready XR applications. In alignment with this objective, XR4ED encompasses key components, such as an open marketplace for XR educational tools, an Ethics Observatory to ensure compliance with EU privacy and inclusiveness standards, and a sustainable ecosystem that fosters collaboration across diverse sectors. The project is coordinated by CYENS Centre of Excellence [L2] in Cyprus and involves partners from eight European countries. Among the deliverables of XR4ED is the Intelligent XR Assistant, an AI-driven system designed to enhance user interaction with XR learning environments by providing personal-



Figure 1: Intelligent XR Assistants while guiding visitors at the Venetian Walls, in the old city of Nicosia.

ized guidance, adaptive content recommendations, and real-time support for both educators and learners.

Intelligent XR Assistant: Bridging AI and XR

The Intelligent XR Assistant integrates, not only XR but also computer animation and AI to develop highly interactive and immersive learning environments for various application domains, including cultural heritage. As part of this endeavour, an AI-powered guide for Nicosia is introduced, offering an enriched exploration of the city's historic centre, which is enclosed by the Venetian Walls (Figure 1). Built by the Venetians in the 16th century to replace the existing Frankish fortifications, it stands as the city's most significant, recognisable, and well-preserved cultural heritage asset. Moreover, Nicosia's historic centre is home to a wealth of cultural heritage treasures, including historical sites and monuments, as well as significant archaeological, archival, and fine art collections preserved in museums and private institutions.

The XR Assistant features realistic avatars created with Character Creator 4, enhanced by Ready Player Me's animation library [L3] to provide natural body language, including gestures, facial expressions, and subtle movements while speaking and when idle. Communication is facilitated through Azure Cognitive Services [L4], enabling multilingual speech-to-text and text-to-speech capabilities, with viseme generation

ensuring accurate lip-syncing across different languages [3]. At its core, the assistant integrates OpenAI's ChatGPT [L5] API to process user inputs and generate contextually relevant responses, allowing dynamic interactions during presentations. Additionally, the system supports a dual-agent mode, where an AI-driven avatar engages in discussions with users on selected topics, offering information, insightful analysis, and complementary viewpoints or contrasting interpretations based on the context [2]. The avatar is depicted in Figure 1.

Built using Unity game engine (version 2022.3.1f1) and Vuforia Engine 10.25 for XR functionalities, the assistant is compatible with both Magic Leap 2 [L6] AR head-mounted displays and Android devices, ensuring accessibility across various platforms [2]. To enhance immersion, the assistant features a panoramic image mode, enabling the on-screen display of contextual backgrounds relevant to the educational content. This capability is particularly valuable in the case of historical sites and monuments, as it facilitates exploration and engagement with locations or environments tied to specific historical contexts, as shown in Figure 2.

The system's architecture also supports collaborative learning, allowing multiple users to connect, interact, and engage with the content simultaneously, fostering a shared educational experience [2]. Overall, the Intelligent XR Assistant leverages cutting-edge XR and AI technologies to provide a customizable, interactive, and immersive educational platform, aiming to transform traditional learning methodologies.

This work has received funding from the European Union's H2020 Programme under Grant Agreement No 739578 (RISE Project) and the Government of the Republic of Cyprus through the Deputy Ministry of Research, Innovation and Digital Policy. It has also received funding from the European Union under Grant Agreement No 101093159 (XR4ED Project).

Links:

- [L1] <https://xr4ed.eu/>
- [L2] <https://ex.cyens.org.cy/>
- [L3] <https://kwz.me/hAU>
- [L4] <https://azure.microsoft.com/en-us/products/ai-services>
- [L5] <https://chat.openai.com/>
- [L6] <https://www.magicleap.com/magic-leap-2>

References:

- [1] S. Boiano et al., "Chatbots and new audience opportunities for museums and heritage organisations," in *Electronic Visualisation and the Arts (EVA)* 2018, pp. 164–171. <https://doi.org/10.14236/ewic/EVA2018.33>
- [2] F. Liarokapis, V. Milata, and F. Škola, "Extended reality educational system with virtual teacher interaction for enhanced learning," *Multimodal Technologies and Interaction*, vol. 8, no. 9, p. 83, 2024. <https://doi.org/10.3390/mti8090083>
- [3] F. Škola, V. Milata, and F. Liarokapis, "Extended reality intelligent assistant for education," *ERCIM News*, no. 137, 13 May 2024. [ISSN: 0926-4981].

Please contact:

Fotis Liarokapis
CYENS – Centre of Excellence, Nicosia, Cyprus
f.liarokapis@cyens.org.cy




Figure 2: Panoramic view of the Intelligent XR Assistants in the old city of Nicosia. In the background is the Church of Panagia Faneromeni (the Revealed One in Greek), built in 1872 to replace an older ecclesiastical structure that existed on the same site.

CulturAI: Advancing Mixed Reality Art Experiences with Large Language Models

by Dimitrios Koukopoulos (University of Patras, Greece),
Christos A. Fidas (University of Patras, Greece), Marios
Belk (Cognitive UX GmbH, Germany)

An innovative Mixed Reality application, funded under the EU Erasmus+ CREAMS project, combines immersive technologies and large language models to revolutionise the way we experience and interact with virtual art exhibitions.

In recent years, the cultural heritage (CH) domain has shown growing interest in exploiting advanced AI-driven tools to enrich user engagement and provide personalized user experiences. This is driven by the increasing interest of museums, art galleries, and other cultural institutions in attracting diverse audiences and offering deeper learning opportunities through digital innovations. However, research into the utilisation of Large Language Models (LLMs) in Mixed Reality (MR) cultural settings remains relatively underexplored.

CulturAI addresses this gap by presenting an MR application that displays holographic artworks while integrating a conversational cognitive assistant. Users can directly ask questions about an artwork's creator, its historical context, or the stylistic details that make it unique. The cognitive assistant responds based on LLM-generated information, delivering immersive storytelling and interactive learning. Figure 1 illustrates how digital artworks are displayed, along with textual responses provided by the cognitive assistant in response to user queries.

This dialogue-based approach can open up new possibilities for visitor engagement, extending beyond merely viewing an exhibition to actively engaging in conversation with an AI-driven virtual guide. To achieve this functionality, CulturAI employs a client-server architecture. On the client side, the MR application, built for the Microsoft HoloLens device, visualizes digital artworks within virtual exhibitions and interprets the user's voice queries. The user simply gazes at a holographic painting, for instance, and utters questions like, "Who was the artist behind this artwork?" or "Which historical event inspired this painting?" The system's user interface is designed to be intuitive and unobtrusive, allowing users to focus on the artwork rather than the technology and device controls. On the server side, a Python-based web service hosts the digital artwork content, processes voice recordings, and uses an LLM to generate responses in near real-time. These responses are then transmitted back to the MR application and displayed as augmented holographic text, providing context-aware insights into artists, historical periods, and stylistic nuances. As a result, visitors can learn about an artist's background, cultural influences, or even the story behind the artwork's creation without pausing their immersive experience.

A preliminary evaluation study was conducted with 39 volunteers who participated and interacted with virtual art exhibitions within an ecologically valid mixed reality context. The study aimed to assess user satisfaction, overall experience, and the perceived trustworthiness of an AI-driven virtual guide. Participants were split into groups, comparing a conventional MR tour (i.e., static text descriptions of artworks) with the interactive, voice-based LLM approach. Analysis of results provided initial evidence of the positive aspect of integrating LLMs in MR applications in terms of user experience and perceived trust. Many participants noted that the responses from the cognitive assistant deepened their engagement, as they could pose spontaneous follow-up questions and receive compelling, context-specific explanations.

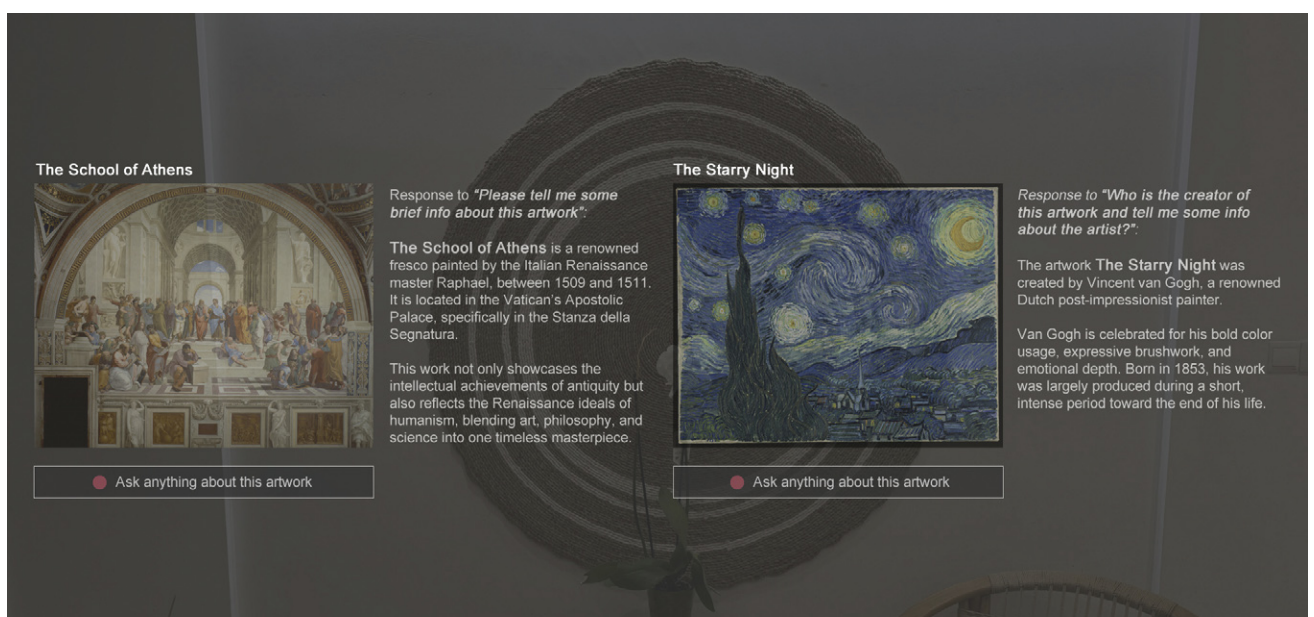


Figure 1: CulturAI application displaying digital artworks along with textual responses provided by the cognitive assistant in response to user queries. Left image depicts "The School of Athens" by Raffaello Sanzio da Urbino (public domain via Wikimedia Commons); Right image depicts "The Starry Night" by Vincent van Gogh (public domain via Wikimedia Commons).

This research work is the outcome of a close collaboration between the University of Patras (Greece) and Cognitive UX GmbH (Germany), combining expertise in Mixed Reality development, AI, and user experience design. The research was partially funded by the Erasmus+ CREAMS project (Project ID: KA220-HED-E06518FA) under the call KA220-HED - Cooperation partnerships in higher education (Greek State Scholarships Foundation (IKY) [L1]).

This research was published as a demo paper [1] at the ACM User Modeling, Adaptation and Personalization (ACM UMAP) conference 2024, where it received the Best Demo Award. The positive reception underscores the project's innovation and its potential impact on interactive experiences in extended reality contexts.

Looking ahead, future work entails integrating more artwork collections and tailoring interactive features for museums and heritage sites worldwide, thereby making CulturAI a more comprehensive platform. On the technical side, the project team envisions improving speech recognition accuracy and exploring advanced user modelling to deliver personalised and adaptive interactions. Another research direction involves investigating the system's effectiveness in enhancing user engagement and retention over extended periods, potentially informing museums about how best to structure digital tours for different demographics. The team also plans to refine the suggested LLM-driven approach to create dynamic narratives that adapt to individual user models by considering the unique characteristics, preferences, and interests of the end-users. For instance, a history enthusiast might receive more detailed political context, whereas an art student might be guided through stylistic analysis and brushwork techniques.

By bridging Mixed Reality displays, real-time LLM-driven dialogue, and cultural storytelling, CulturAI demonstrates how technological innovation can expand access to the arts and transform learning experiences in museum and heritage contexts.

Link:

[L1] <https://creams-project.eu>

Reference:

[1] N. Constantinides, et al., "CulturAI: Exploring Mixed Reality Art Exhibitions with Large Language Models for Personalized Immersive Experiences", in Adjunct Proc. of the 32nd ACM Conf. on User Modeling, Adaptation and Personalization (UMAP Adjunct '24), ACM, 102–105, 2024. <https://doi.org/10.1145/3631700.3664874>

Please contact:

Dimitrios Koukopoulos
University of Patras, Greece
dkoukopoulos@upatras.gr

Christos A. Fidas
University of Patras, Greece
fidas@upatras.gr

Towards more Inclusive Museums Using AI: Guidance in Sign Language

by Dimitrios Kosmopoulos (University of Patras), Evanthia Papadopoulou (Archaeological Museum of Thessaloniki), Antonis Argyros (FORTH-ICS)

The SignGuide project has developed an interactive museum guide system for deaf and hard-of-hearing visitors using mobile devices to promote inclusion for this social group. The system is capable of understanding visitors' questions in sign language and provides additional content also in sign language using AI methods. The system has been deployed in the Archaeological Museum of Thessaloniki.

SignGuide: Making Museums Accessible for Deaf and Hard-of-Hearing Visitors

Museums are designed to educate, inspire, and preserve cultural heritage, yet for deaf and hard-of-hearing (DHH) visitors, they often remain inaccessible. Traditional museum experiences heavily rely on text-based descriptions and audio guides, which can pose challenges for those who primarily communicate in sign language rather than written language. Research has shown that many deaf individuals struggle with written content because sign language has a different grammatical structure from spoken languages, making direct translations inadequate for full comprehension [1]. As a result, DHH visitors often miss out on the depth of information available to hearing visitors, leading to an unequal experience.

The SignGuide project addresses these accessibility barriers by providing an interactive guidance system through a mobile application. Unlike conventional accessibility solutions that rely on captions or text summaries, SignGuide prioritizes sign language as the primary mode of communication, offering sign language videos and avatars to enhance engagement with museum exhibitions. This approach ensures that DHH visitors can explore exhibits autonomously, ask questions in sign language, and receive responses in their preferred communication mode.

A Real-World Use Case

Maria, a deaf visitor, enters the Archaeological Museum of Thessaloniki and notices a digital display announcing the availability of sign language guidance. She downloads the SignGuide app and is welcomed by a sign language introduction explaining how to use the app.

As she walks to the first exhibit, she notices a SignGuide logo on the display. She opens the app, points her camera at the exhibit, and within seconds, the app recognizes the artifact using augmented reality and displays a sign language video explaining its historical background [2].

Curious to learn more, Maria taps on the "Ask a Question" button. She selects a pre-recorded question in sign language asking about the exhibit's cultural significance. The app instantly retrieves a video response in sign language, providing



Figure 1: As soon as the app visually recognizes an exhibit, it can respond to user queries in sign language and retrieve related content in the same format.

her with more detailed insights (see Figure 1). Later, Maria visits another exhibit and chooses to record her own question using the front camera. The system processes her query, finds a matching pre-recorded answer, and delivers it in sign language. For the first time, Maria feels truly engaged in a museum experience, enjoying an equal level of access to information as hearing visitors.

The consortium

SignGuide was funded through the Research Create Innovate action in Greece (2020-2023), T2EDK-00982, by the consortium of the University of Patras (coordinator, sign language understanding, sign language linguistics), Foundation for Research and Technology - Hellas (human motion tracking, avatars), Bioassist SA (system integration) and the Archaeological Museum of Thessaloniki, which hosts the exhibition. The application is available on Apple Store and on Google Play [L1].

Interactive Q&A in Sign Language: One of the most innovative features of SignGuide is its interactive Q&A system, which allows visitors to ask questions in sign language and receive responses in the same format. Visitors can either:

- Select from a list of pre-recorded questions in sign language
- Record their own question using the phone's front-facing camera.

Once recorded, the app processes the visitor's query using AI-powered sign language recognition. The system tracks hand movements using a skeleton-tracking algorithm, extracts hand pose sequences, eventually extracts visual summaries and then feeds them into a Transformer-based neural network. This deep learning model matches the recorded sign language query to a set of known questions. After confirming the correct translation, the app retrieves a pre-recorded video response or generates an answer using a sign language avatar. Currently, the system supports over 300 possible questions for the most popular exhibits in the Archaeological Museum of Thessaloniki.

Impact and Reception: The SignGuide project has been warmly received by the DHH community and has been de-

signed after careful consideration of their needs. Many schools for the Deaf have visited the Archaeological Museum of Thessaloniki specifically because of this accessibility feature. For many students, it was their first fully accessible museum experience, highlighting the importance of inclusive cultural engagement. By bridging the communication gap between museums and DHH visitors, SignGuide not only enhances accessibility but also promotes a deeper appreciation of cultural heritage among underserved communities.

Future Plans: We plan to significantly upgrade the provided services by: (a) leveraging large language models and avatars to synthesize dynamically the retrieved content in sign language, (b) keeping up with the advancements in video and sequential data processing, (c) providing information before, as well as after the visit, including educational content and games, and (d) involving more sign languages – currently the Greek Sign Language is supported.

Link:

[L1] <https://www.signguide.gr>

References:

- [1] M. Nikolarazi, I. Vekiri, and S. Easterbrooks, "Investigating deaf students' use of visual multimedia resources in reading comprehension," *Am. Ann. Deaf*, vol. 158, no. 4, pp. 363–377, 2013.
- [2] D. Koulouris, et al., "Utilizing AR and hybrid cloud-edge platforms for improving accessibility in exhibition areas," in *Proc. IFIP Int. Conf. Artif. Intell. Appl. Innov.*, Springer, 2023, pp. 123–134.
- [3] E. G. Sartinis, E. Z. Psarakis, and D. A. Kosmopoulos, "A 3D wrist motion-based sign language video summarization technique," *Pattern Recognit. Lett.*, vol. 195, pp. 45–53, 2025.

Please contact:

Dimitrios Kosmopoulos
University of Patras, Greece
dkosmo@upatras.gr

TCRMQ: An Application for Generating CIDOC-CRM SPARQL Queries from Text by Using Large Language Models

by Michalis Mountantonakis and Yannis Tzitzikas (FORTH-ICS and University of Crete)

TCRMQ (Text-2-CIDOC-CRM Query) is a research prototype that can generate SPARQL queries over CIDOC-CRM Knowledge Graphs using a novel two-stage ontology path pattern method and Large Language Models (LLMs). The generated SPARQL queries can include in the filtering conditions information that are not even part of the KG, by exploiting knowledge from the LLM.

There is a strong need for exploiting Artificial Intelligence (AI) techniques, to enhance the accessibility and reusability of Knowledge Graphs (KGs) of Cultural Heritage domain, which have been modelled through the CIDOC-CRM ISO Standard [1]. Hence, a relevant task is to enable the answering of natural questions over CIDOC-CRM based KGs, by generating SPARQL queries from a given natural question. This is a highly challenging task, especially for sophisticated (event-based) ontologies like CIDOC-CRM, since in most cases one has to derive queries with long property path expressions to answer a given question [2]. To tackle this difficulty, we present the research prototype TCRMQ (i.e., Text-2-CIDOC-CRM Query) [L1], which is an application that can generate SPARQL queries over CIDOC-CRM KGs, by exploiting a novel two-stage ontology path pattern method that combines

knowledge from a given KG and from Large Language Models (LLMs) [2].

Regarding the functionality, as we can see in Figure 1, TCRMQ receives a question from the user in natural language, such as “Give me the number of artworks for each artist born after the end of the first World War” and then the user selects a method and the desired KG to answer the given question. The application uses a novel two-stage method combining Ontology Path Patterns and Knowledge from LLMs [2] for generating the SPARQL query. First, the method sends a prompt with candidate ontology path patterns and receives a prediction with the most relevant path patterns for answering the question. Then, a second prompt is sent to the LLM, for generating the SPARQL query for the given question, based on the predictions. For instance, in Fig. 1, TCRMQ managed to predict the long property path that connects an artwork with the corresponding artist and his/her birth date, and then it generated the corresponding multi-hop SPARQL query. Subsequently, the generated SPARQL query (see the bottom-left part of Figure 1) is sent to the SPARQL endpoint of the desired KG and the application retrieves the answer and presents the results of the query to the user (see the right part of Figure 1).

A very interesting finding regarding the synergy between LLMs and KGs is that LLMs “know” background knowledge (dates, names, etc.) that is not present in the KG, and this synergy can yield queries with the right filtering conditions, even if such information is not present in the KG. For example, in Figure 1, the KG did not contain any information about the dates of the First World War. However, LLMs (such as ChatGPT) include this relevant information and, in our case, the LLM added the desired date in the filtering condition of the SPARQL Query (i.e., “1918-11-11”). Another functionality of the TCRMQ is that the user can (i) compare several methods (based on different prompts) for generating the SPARQL

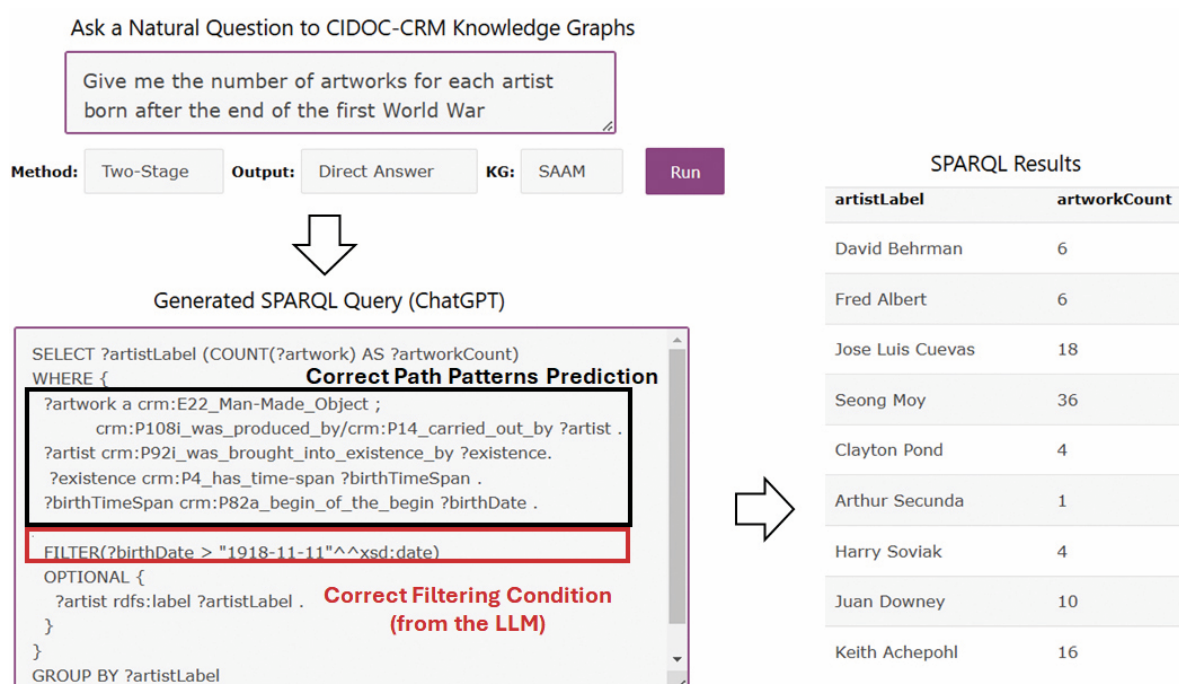


Figure 1: The functionality of the TCRMQ research prototype.

query, (ii) use different ChatGPT versions (currently it supports versions 3.5, 4-mini and 4), (iii) ask a question in languages other than English, and iv) edit the generated SPARQL query, e.g., to restrict the results and/or to correct possible errors in the generated query.

Concerning evaluation results, we have tested the methods of TCRMQ by using a benchmark [L2] with 100 questions over two real CIDOC-CRM based KGs including data about artworks. The benchmark consists of five categories of questions, i.e., each category includes 20 questions requiring a path with length a) 1, b) 2, c) 3, d) 4 and e) mixed paths (combining paths of different length). The evaluation was performed by using the ChatGPT version 3.5. Regarding the results, we managed to generate the correct SPARQL query for 66 out of the 100 questions [2], by using the two-stage method, whereas the corresponding accuracy for the baseline method was 19 out of 100 questions. As regards the different categories of questions, the accuracy was lower for questions requiring following either a path of length 4 or a mixed path, compared to the other categories. In the most erroneous cases, either a totally incorrect path or a longer path was selected. All the details for the benchmark and the evaluation can be found in the paper [2] and on the GitHub page [L2].

As future work, we plan to evaluate the methods by also using newer versions of ChatGPT, by also expressing the questions in other languages, to add more KGs in an automatic way (including also other KGs using event-based models) and to perform an analysis of the path patterns of all the available CIDOC-CRM KGs.

Links:

[L1] <https://demos.isl.ics.forth.gr/Text2CIDOC/>

[L2] <https://github.com/mountanton/CIDOC-QA-using-LLMs>

References:

- [1] Y. Tzitzikas, et al., “CIDOC-CRM and machine learning: a survey and future research”, *Heritage*, 5(3), pp.1612-1636, 2022.
- [2] M. Mountantonakis and Y. Tzitzikas, “Generating SPARQL Queries over CIDOC-CRM using a Two-Stage Ontology Path Patterns Method in LLM Prompts”, *ACM Journal on Computing and Cultural Heritage*, 18, 1, Article 12, March 2025, 20 pages.
<https://doi.org/10.1145/3708326>

Please contact:

Michalis Mountantonakis
FORTH-ICS and University of Crete, Greece
mountant@ics.forth.gr

Yannis Tzitzikas
FORTH-ICS and University of Crete, Greece
tzitzik@ics.forth.gr

AI Multilingual Search Platform for EU Audiovisual Media Archives

by Pilar Orero, Chiara Gunella, and Sarah McDonagh
(Universitat Autònoma de Barcelona)

Culture and linguistic diversity in Europe is both a jewel and a barrier to communication. Broadcaster's archives are part of the EU cultural heritage –inaccessible for many reasons. The MOSAIC project aims to develop tools for multilingual translation, automatic subtitling, and AI-driven content adaptation. MOSAIC seeks to empower broadcasters and media producers by making their content available to a wider audience, enhancing cultural exchange and unity across Europe.


European broadcasters, news agencies, and corporations are known for high-quality media content and information production, distribution, and consumption and for the richness of their content archives, containing a wealth of cultural, political, historical, and artistic content in various formats, including film/video, television, radio, and digital media. Nevertheless, digital platforms outside Europe already exert control over the media landscape, leading to a fragmentation of accessible knowledge within the European Union. The potential of Europe's richness remains underused with a serious risk that Europe will lose valuable resources and archival material. In this scenario, the adoption of artificial intelligence (AI) technologies by organisations, and thus the creation of media content using advanced techniques, is essential not only for the cultural sovereignty of Europe but also for the preservation and presentation of heritage, education, and entertainment across a European context. Creating a more unified European media identity, based on a fragmented but common cultural heritage, is a major challenge: on one side, both public and private broadcasters have to ensure that national regulatory authorities are independent of political or commercial influence to ensure media freedom and pluralism, on the other side, the common challenge is the constant need to adapt to new technologies to stay relevant and to have the chance to make the European media market competitive worldwide and European media content archives an unparalleled source of cultural richness.

Artificial Intelligence (AI), Natural Language Processing (NLP), Natural Language Understanding (NLU), Language Technologies (LTs), and Speech Technologies (STs) have the potential to enable multilingualism technologically but, according to the META-NET White Paper Series “Europe's Languages in the Digital Age” [1] published in 2012, our languages suffer from an extreme imbalance in terms of technological support: English is well supported through technologies, tools, and datasets, but languages such as Maltese, Estonian or Icelandic still have very poor support.


The goal is to enable multilingualism technologically since “the EU and its institutions have a duty to enhance, promote and uphold linguistic diversity in Europe” (European Parliament 2018) [2]. Today, Generative AI (GenAI) can cre-

Scenarios

Scenario 1:
An organization prepares enhanced content for internal use by media professionals from the same organization



Scenario 2:
An organization prepares enhanced content and distributes it to consumers



Scenario 3:
An organization requests content from another organization and enhances it for use by media professionals




Figure 1: MOSAIC scenarios.

ate all kinds of media, including texts, sounds, videos and 3D content. Its use by news organisations and broadcasters is still in its infancy. It has the potential to reshape the concept of creation and affect the operation and business models across the media and cultural sectors, with the danger of fake news being created at scale. It is expected that the impact on the media and cultural and creative industries (CCI) will be significant, but strong research and innovation support is needed to fully benefit from these new opportunities [L1].

The GenAI market is expected to grow substantially in the next few years. For example, a recent report from Sopra Steria shows that it could go from around 8 billion USD in 2023 to more than 100 billion by 2028. At the same time AI can also generate intelligent media services that adapt to user needs, including environmental factors and personal capabilities, thereby providing accessible multi-language, multimodal media services as prerequisite to enjoying any XR media content. However, the full potential of AI still needs to be further exploited within ethical and legal boundaries, ethical guidelines and regulations should be developed to address the issues related to such models such as sustainability, safety, intellectual property, bias, explainable AI, and trustworthiness.

In terms of AI for Media Access Services, NEM SRIA 2024 [L1] identified some major challenges like i) streamlining the circulation of audiovisual (or video) programs through machine translation, while humans focus on the quality of work; ii) encouraging synergies and convergence between subtitling and the development of multilingualism or the integration of foreigners (e.g., migrants); iii) developing AI tools for automatic translation from speech to subtitles, from text to Sign Language, and from Sign Language to text; iv) Develop AI tools for robust automatic translation of subtitles (multi-languages). In this context, MOSAIC [L2] aims to develop a prototype European AI-enhanced platform, serving

as a central, scalable hub for broadcasters and news creators, distributors and consumers. The platform, as can be seen in Figure 1, will leverage knowledge repositories by a sophisticated, multilingual and multimodal AI-based integrated system that links to producers and harnesses the abundance and richness of cultural heritage, media and news repositories, while providing a source of monetisation from accumulated knowledge repositories.

MOSAIC 2024-2026 is co-funded by the DIGITAL EUROPE program under grant no. 479833.

Links:

[L1] <https://kwz.me/hGC>

[L2] <https://mosaic-media.eu/>

References:

- [1] G. Rehm and H. Uszkoreit, "Language Technology Support for Polish," in *The Polish Language in the Digital Age*, Springer, 2012, pp. 52-67.
- [2] G. Rehm, A. Way (eds.), "European Language Equality", Springer, 2023. <https://doi.org/10.1007/978-3-031-28819-7>

Please contact:

Pilar Orero

Universitat Autònoma de Barcelona, Spain

pilar.orero@uab.cat

Taming AI both as Tool and Artefact for Computer Museums

by Christophe Ponsard (University of Namur), Caroline Djambian, and Emrick Poncet (University of Grenoble)

The current AI revolution is shaking up the cultural world and its creation process. In computer museums AI can be seen both as a powerful curatorial tool and an artefact requiring careful processing. We present and illustrate here the dual roles of AI in enhancing museum preservation and storytelling by providing a few specific scenarios, while being part of the story itself and raising concerns about cultural authenticity and ethics.

Artificial Intelligence (AI) has a long history prior to its current boom resulting from the combination of deep learning algorithms, powerful hardware, and huge training datasets. It is increasingly being adopted by curators as a tool to operate their museums with several reported uses [1]. Interestingly, for computer museums, AI is also worth being analysed and presented as part of the computer storyline itself, especially as the leadership in evolution appears to be shifting from humanity to AI, raising profound ethical questions about an entity devoid of consciousness yet driven by financial interests.

From a historical perspective, a typical museum presentation could be Figure 1, revealing that the current revolution, seen as disruptive, is the result of successive innovation cycles over

several decades. It can also be related to other enabling revolutions such as micro-electronics, big data and learning algorithms. More specific views could also focus on key people, on the move from circles of experts to the general public, and on environmental and societal impacts. AI is also strongly reshaping culture by its ability to generate text, pictures, music, and videos triggering debates about creativity, biases, deepfakes, meddling and even humanity.

Many missions and daily activities of a museum can be supported by AI tools:

- For preservation: support inventory efforts by aiding artefact identification through image recognition, leveraging the public web and shared museum collections. Use text generation models (LLM) fed with technical databases to produce documentation drafts.
- For research and exhibition, create alternative thematic pathways for exploring exhibitions, offering diverse perspectives as outlined above.
- For education: use text and image generative AI to develop material adapted for school or family events, including gamification activities, also achieving multilingual support.

Of course, the use of AI should be prepared and supervised. As a detailed use case, let's consider the main exhibition of the NAM-IP museum located in Belgium [L2]. The full visit follows a raw timeline of about 700 events covering technological discoveries, computing machines, innovations, failure stories, and general context. Currently, only a few viewpoints are proposed to the public as it is hard work. However, generative AI can produce them in a few seconds to suggest a museum tour focusing on a given technology (e.g. punch card machine), country (e.g. UK) or topic (e.g. security, AI). Raw text

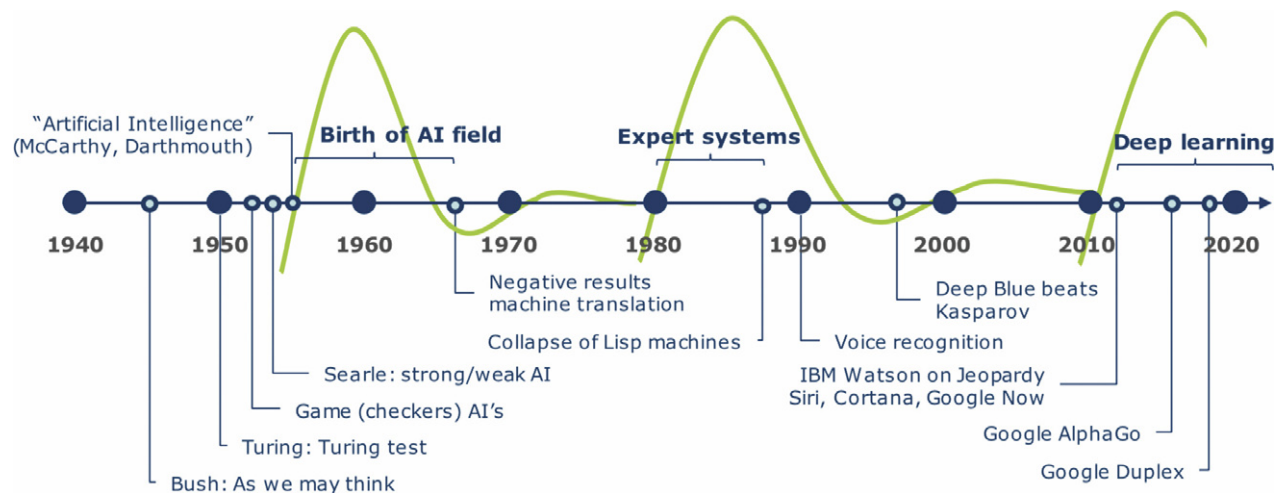


Figure 1: AI timeline (source [L1]).




Figure 2: Automatically generated interactive timeline.

output can be transformed in static or dynamic visualization, e.g. with vis.js [L3] as shown in Figure 2 for punched card technology.

However special care is required: if not correctly prompted, the chatbot might inject information from uncontrolled sources or even hallucinate. It can introduce bias but can also help the curator to detect potential existing ones (e.g. in our case women were underrepresented in the existing timeline). This means the process should be guided and cross checked by human experts.

To learn to address such issues, we are pursuing our work through the ITinHeritage research project, backed by an international network of science or computing museums. Together, we explore the possibilities offered by AI aided by open data, semantic web, and Virtual Reality (VR) [L4]. So far, our collaboration has allowed us to aggregate European computing cultural heritage metadata into a Knowledge Graph and make it available as Linked Open Data (LOD). By structuring it with a diachronic domain ontology, we could enhance its meaning and improve the quality of AI processing [2]. On this basis, we also aim to apply AI techniques on extended corpora, to generate historical models of computer science evolution, supporting the work of historians and museums. This will give the opportunity to compare the completion of LLMs with ontologies constructed by experts, and to observe its biases with regard to historical facts that can be verified by experts with the help of ontology and timeline navigation operators [3]. We will then attempt to derive forward-looking visions of technological developments. Our aim is to use AI to position museums as actors of a science in the making, and not just as collectors of memory. The results produced by the AI could be made available to the public on the web, in VR or in museums. In this way, we want to provide museums with significant assistance in managing and mediating their collections using these news tools.

References:

- [1] M. Münster et al., “Artificial Intelligence for Digital Heritage Innovation: Setting up a R&D Agenda for Europe,” *Heritage*, vol. 7, pp. 794–816, 2024.
- [2] C. Djambian et al., “New terminological approaches for new heritages and corpora: The ITinHeritage project,” in *Proc. of MDTT*, Grenade, Spain, 2024.
- [3] C. Ponsard, A. Masson, and W. Desmet, “Historical Knowledge Modelling and Analysis through Ontologies and Timeline Extraction Operators: Application to Computing Heritage,” in *Proc. of MODELSWARD*, 2022.

Links

- [L1] <https://kwz.me/hGh>
- [L2] <https://www.nam-ip.be>
- [L3] <https://visjs.org>
- [L4] <https://itinheritage.wordpress.com>

Please contact:

Christophe Ponsard
NAM-IP, University of Namur, Belgium
christophe.ponsard@unamur.be

Caroline Djambian, University of Grenoble, France
caroline.djambian@univ-grenoble-alpes.fr

Digital Transformation in Legal History Through Automatic Machine Learning Annotation

by Giovanni Puccetti (CNR-ISTI), Vincenzo Roberto Imperia (University of Palermo), and Andrea Esuli (CNR-ISTI)

An interdisciplinary project leverages AI to automatically annotate legal references in the Liber Extra's Ordinary Gloss, creating a valuable index for researchers in Legal History and Religious Studies. Using efficient machine learning techniques, we achieved near-perfect accuracy in automatically identifying and linking thousands of legal references. The resulting open-source workflow offers a replicable solution for similar tasks in the humanities.

The study of medieval legal texts, such as the Liber Extra's Ordinary Gloss [1], is crucial for understanding the evolution of legal and religious thought. These texts contain “allegations”, legal references, i.e., link to a norm that support the interpretation of another norm, that form the backbone of medieval legal argumentation. Manually annotating these references is a labor-intensive process, requiring highly specialised domain expertise and significant time. In this context we proved that machine learning can provide a scalable and efficient solution, enabling the automatic extraction and annotation of these references with high accuracy.

This project [3], part of the ITSERR initiative [L1], developed an automatic annotation system for legal references in the Liber Extra's Ordinary Gloss. The Liber Extra, a cornerstone of medieval canon law, is accompanied by the Glossa Ordinaria, a commentary that includes thousands of legal references. These references are essential for understanding the interpretative framework of the text but are challenging to navigate due to their volume and complexity. By combining expert manual annotation with machine learning, we created an index of 41,784 legal references, linking each to specific lemmas, chapters, and titles in the Liber Extra (see Figure 1).

The process began with a domain expert annotating a subset of the text using INCEpTION [2], a collaborative annotation platform designed for efficient text annotation. The expert annotated 12 out of 185 titles, resulting in 4,578 annotations. These annotations served as the training data for a machine learning model, which was then applied to the entire text.

We tested two machine learning approaches: Conditional Random Fields (CRF), a traditional statistical method, and transformer-based models such as BERT and Latin BERT. CRF models, which excel at sequence labelling tasks, achieved 97.8% accuracy in identifying legal references, outperforming the transformer-based models. The CRF model's success can be attributed to its ability to capture the contextual dependencies in the text, such as the specific abbreviations and numbering conventions used in medieval legal citations. In contrast, while Latin BERT performed better than the original

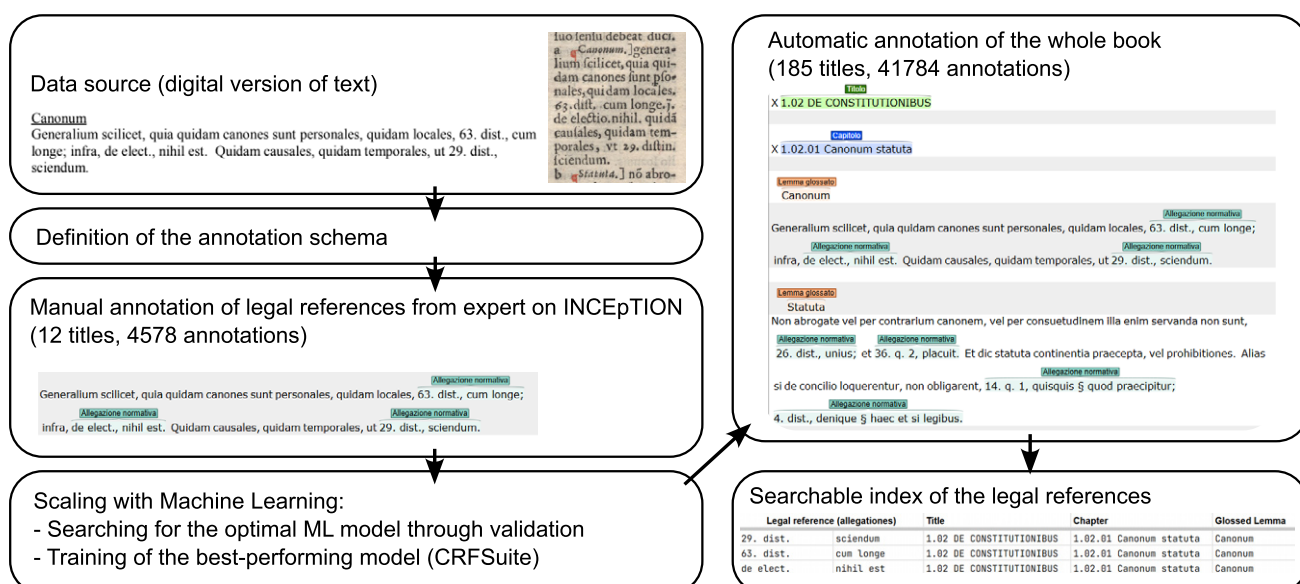


Figure 1: An illustration of the process that led to construction of the index of all the legal references in the Liber Extra's Glossa Ordinaria.

BERT due to its specialisation in Latin, it still fell short (-5.4%) of the CRF model's accuracy.

The CRF model's efficiency was another advantage. CRF uses the CPU of a desktop computer for training, making it accessible to small research teams with limited computational resources. In contrast, training transformer-based models demanded significant – and much more costly – GPU resources, highlighting the practicality of CRF for similar projects in the humanities. Another term of comparison in favor of CRF is the size of the trained model, which is just 1 MB for CRF compared to the 423 MB of Latin BERT.

Once trained, the CRF model was used to annotate the entire Liber Extra's Ordinary Gloss, identifying 41,784 legal references. Each reference was linked to a specific lemma, chapter, and title in the Liber Extra, creating a comprehensive index, which was further validated by the domain expert, showing a very high accuracy [L2]. The index not only facilitates navigation but also enables new forms of analysis. For example, researchers can now compute the frequency of use in references of any norm, their distribution in text, supporting the exploration of their role and significance in medieval legal thought. The CRF model also found a few cases of non-legal references, such as citations to Gospel passages, highlighting the complexity of medieval legal texts.

The index is thus a valuable resource for researchers in Legal History and Religious Studies, enabling the exploration of connections between legal norms and their interpretations, shedding light on the intellectual context of medieval jurists.

This project highlights the transformative potential of AI in the study of historical texts, providing tools that make vast corpora accessible and analysable. The creation of coordinated and interconnected databases of normative texts and their commentaries could significantly impact the field, enabling large-scale comparative studies and the development of advanced retrieval systems.

By combining domain expertise with machine learning, we created a resource that not only supports current research but also paves the way for future discoveries. The index and code [L3] are publicly available, to be used as tools and as guidelines for similar projects in the humanities.

The ITSERR project is funded by the European Union - NextGenerationEU (CUPB53C22001770006).

Links:

[L1] <https://www.itserr.it/>

[L2] <https://doi.org/10.5281/zenodo.14381709>

[L3] https://github.com/aesuli/CIC_annotation

References:

- [1] E. Reno, "The Digital Decretals", 2024.
- [2] J.-C. Klie et al., "The INCEpTION Platform: Machine-Assisted and Knowledge-Oriented Interactive Annotation," 2018.
- [3] A. Esuli et al., "Automatic Annotation of Legal References (Allegations) in the Liber Extra's Ordinary Gloss", IRCDL 2024

Please contact:

Giovanni Puccetti
CNR-ISTI, Pisa, Italy
giovanni.puccetti@isti.cnr.it

AI-driven Extraction of Structured Data from a Large Newspaper Archive

by Nikos Kontonasios, Yannis Tzitzikas (FORTH-ICS and University of Crete), and Pavlos Fafalios (FORTH-ICS and Technical University of Crete)

The project PortAda investigates ship arrivals at the ports of Barcelona, Marseille, Havana, and Buenos Aires from 1850 to 1914, aiming to generate an open-access database containing thousands of records related to 19th-century maritime trade. As part of this effort, the Centre for Cultural Informatics of FORTH-ICS is developing a data extraction pipeline for ship arrivals recorded in the historical newspaper *Le Sémaphore de Marseille* (1827–1944). The pipeline integrates document layout analysis, optical character recognition (OCR), and generative AI-driven techniques to extract relevant data from scanned newspaper images and store it in a well-structured, machine-readable format for further analysis.

Historical newspaper archives serve as invaluable sources of information, documenting not only political and economic changes but also cultural, social and technological developments [1, 2]. PortADA is a European project (MSCA Staff Exchanges) in the field of maritime history that draws on historical maritime newspapers to study maritime trade, economic history, and port activities in the 19th and early 20th centuries [L1]. It focuses on ship arrivals at the ports of Barcelona, Marseille, Havana and Buenos Aires from 1850 to 1914. The project's objective is to extract ship arrivals data from local newspaper archives and develop an open-access database containing thousands of records that document 19th-century maritime trade.

The project is coordinated by the University of Barcelona and involves multiple partners, including the Institute for Mediterranean Studies of FORTH (IMS-FORTH), the Institute of Computer Science of FORTH (ICS-FORTH), the Autonomous University of Madrid, the Maritime Museum of Barcelona, and institutions in Argentina and Cuba. IMS and ICS-FORTH play a key role in extracting ship arrival records for the port of Marseille, using the historical newspaper *Le Sémaphore de Marseille* as its primary source, which consistently reported all ship arrivals in Marseille in its daily issues. The newspaper's archives span from 19 December 1827 to 19 August 1944, with 35,703 issues published during this period. The full collection is publicly accessible through the Bibliothèque Nationale de France (BNF)'s dedicated newspaper archive website [L2].

Given the scanned images of a newspaper issue, the task is to extract well-structured data located in a specific newspaper section. In the case of the *Le Sémaphore de Marseille*, the focus is on ship arrivals, which are documented in a section named "ARRIVÉES", as shown in Figure 1. In this section, information about each ship arrival is presented in small paragraphs arranged vertically on the page. A ship arrival paragraph contains the following information: port of origin, departure date, port and date of call, ship type, ship's flag, ship name, ship tonnage, name of captain, cargo information, consignee of cargo, consignee of ship, intelligence (events or interactions during the journey).

Automating the extraction of this data from all newspaper issues presents several interconnected challenges. The first challenge is segmenting complex newspaper layouts into distinct text blocks while preserving the correct reading order. The second challenge involves accurately locating and isolating the ship arrivals section within each newspaper issue. The third challenge is parsing this section into individual paragraphs, each corresponding to a single ship arrival. The fourth challenge is converting the unstructured text of each paragraph into a structured format with specific data points. Finally, the fifth challenge stems from the poor quality of historical newspaper images, which suffer from physical deterioration, fading, smudging, skew, and curvature, all of which can significantly impact text recognition [3].

To address these challenges, we have developed an automated pipeline of integrated solutions that transforms the scanned newspaper images into a structured dataset. The various stages of this process are illustrated in Figure 2 and described in detail below.

Text Block Identification: The first stage establishes the layout and structure of the newspaper page by separating text into individual, clearly defined blocks. The tool used for this task is Arcanum's newspaper segmentation API [L3]. It takes as input an image of a newspaper page and outputs a JSON file containing information about the distinct blocks.



Figure 1. Example of a newspaper section providing information about ship arrivals (source of newspaper image: retronews.fr).

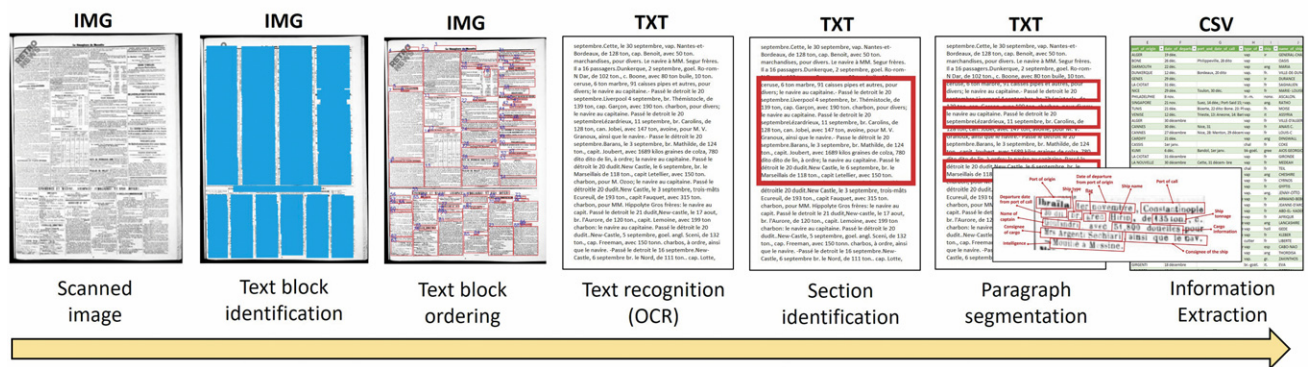


Figure 2. Overview of the information extraction process.

Text Block Ordering: The text blocks identified in the previous step are often not in a sequential reading order due to the layout complexities. To address this, a dedicated software application was developed, which first determines the segments contained in the newspaper page by applying a colour filling to all text blocks using OpenCV [L4]. The blocks are then horizontally stretched so that all blocks within the same segment merge together, effectively highlighting the distinct segments. Then, the columns within each segment are identified using a similar method. Finally, each segment and its respective columns are sorted to ensure a coherent reading order.

Text Recognition (OCR): Once ordered, the text blocks are processed using Google Vision AI to convert the visual text in the images into machine-readable format (text data). Despite its high accuracy, the OCR output can still contain errors, often due to image quality issues.

Section Identification: Because of possible OCR errors, identifying the arrivals section requires more than a simple exact match search for “ARRIVÉES”. To address this, the Levenshtein distance method has been applied, which locates the word closest to “ARRIVÉES” as well as a word resembling “DEPARTS” that appears after it and denotes the end of the arrivals section.

Paragraph Segmentation: The segmentation of the ship arrivals section into individual paragraphs is carried out using ChatGPT 4o, leveraging its text processing capabilities to identify paragraph breaks and formatting cues within the text. For this, a carefully designed prompt was employed which directs the AI model to segment the text into paragraphs without modifying the content.

Information Extraction: The distinct data elements for each ship arrival are extracted from the segmented paragraphs by using ChatGPT's natural language processing capabilities. A detailed prompt was used to guide the extraction process, which explicitly defines the fields to be extracted and instructs the AI model to focus on the raw input text without attempting to interpret or correct OCR errors, thereby preserving the authenticity of the data.

Preliminary evaluation results demonstrate high pipeline performance, achieving robust accuracy in text block ordering, acceptable OCR quality with an average accuracy of 82.3%, and high precision in paragraph segmentation (95.7% F1-

score). The information extraction stage performed reliably, achieving an average F1-score of 96% across twelve information fields. These results highlight the system's strengths while also identifying areas for future refinement.

A key strength of the pipeline is its modularity, which allows for targeted adjustments and improvements to individual components. This flexibility ensures that the system can be easily adapted to suit other documents or extraction tasks.

This work has received funding from the European Union's Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101129889 (Project PortAda).

Links:

- [L1] <https://www.proyectoportada.eu/>
- [L2] <https://www.retroneurs.fr/titre-de-presse/semaphore-de-marseille>
- [L3] <https://www.arcanum.com/en/newspaper-segmentation/>
- [L4] <https://opencv.org/>

References:

- [1] A. Bingham, “The digitization of newspaper archives: Opportunities and challenges for historians,” *Twentieth Century British History*, vol. 21, no. 2, pp. 225–231, 2010.
- [2] B. Nicholson, “The digital turn: Exploring the methodological possibilities of digital newspaper archives,” *Media History*, vol. 19, no. 1, pp. 59–73, 2013.
- [3] J. Jarlbrink and P. Snickars, “Cultural heritage as digital noise: Nineteenth century newspapers in the digital archive,” *Journal of Documentation*, vol. 73, no. 6, pp. 1228–1243, 2017.

Please contact:

Pavlos Fafalios
FORTH-ICS and Technical University of Crete, Greece
fafalios@ics.forth.gr

Automatic Extraction of Regesta for Medieval Latin Text Summarization

by Giovanni Puccetti (CNR-ISTI), Laura Righi (University of Modena and Reggio Emilia), Ilaria Sabbatini (University of Palermo), and Andrea Esuli (CNR-ISTI)

We produced a novel dataset of 4,533 medieval Latin regesta (summaries) paired with full texts, extracted through a meticulous pipeline involving manual annotation, custom model training, text extraction, and post-processing to ensure high-quality, structured data for AI-driven summarization tasks.

The dataset enables AI-driven summarization of historical documents. We evaluated the performance of Large Language Models (LLMs) in generating Latin regesta, offering insights into the challenges and potential of AI in digital humanities.

The REVERINO dataset (REgesta generation VERSus latIN summarizatiOn) [1], developed under the ITSERR [L1] project, focuses on summarizing medieval Latin documents, particularly pontifical texts from the 13th century. Regesta, concise summaries of historical documents, are essential for scholars studying large collections of Latin documents. REVERINO pairs regesta with their corresponding full texts and apparatus (i.e., bibliographic information on the source of the regestum), extracted from two collections: Epistolae saeculi XIII e regestis pontificum Romanorum selectae (MGH, 1216-1268) and Les Registres de Gregoire IX (Auvray, 1227/41).

Generating the REVERINO Dataset

The input data for the process consist of high-resolution images from the pages of the volumes in the two collections. Creating the dataset involved a four-step pipeline:

1. Annotation: Manual annotation of pages by domain experts using the eScriptorium [2] platform to identify and segment

text areas and lines. In this phase we simplified the work of experts requesting that they identify the main content of the page. The experts were not required to separate the different parts of the main content (i.e., regesta, apparatus, full text), as this was done automatically in post-processing.

2. Training: The annotated data allowed us to train customized segmentation models for the collections. The models segmented the main content, separating it from any additional content (e.g., titles, headers, footers, footnotes, line numbers). Each collection required a specific model, as the MGH collection uses a single-column format, while the Auvray collection uses a more complex two-column format, requiring additional training data to improve segmentation accuracy.
3. Extraction: We used the trained models to segment and perform OCR to extract text from all pages. This step produced a continuous stream of text that includes the regesta, full texts, and apparatus, but without labels to distinguish them.
4. Post-processing: We could define simple heuristics based on content and positional information, leveraging the accurate information regarding text position and size produced by the trained models, to accurately split the text into regesta, apparatus, and full texts, producing the structured information we aimed to obtain.

The final dataset comprises 2,283 pairs from the MGH collection and 2,250 from Auvray. A t-SNE analysis (Figure 2) on the text of regesta revealed distinct clusters for each collection, highlighting their unique characteristics due to differences in topics, sources, and editorial style (t-SNE stands for t-distributed Stochastic Neighbor Embedding, a dimensionality reduction technique commonly used to visualise high-dimensional data in a lower-dimensional space).

Training LLM on Regesta Generation

We used REVERINO as a benchmark for evaluating the ability of LLMs to generate regesta. Three models—GPT-4, Llama 3.1 70b, and Llama 3.1 405b—were tested using two approaches:

1. Direct Summarization (format): The model generates a regestum directly from the Latin full text.

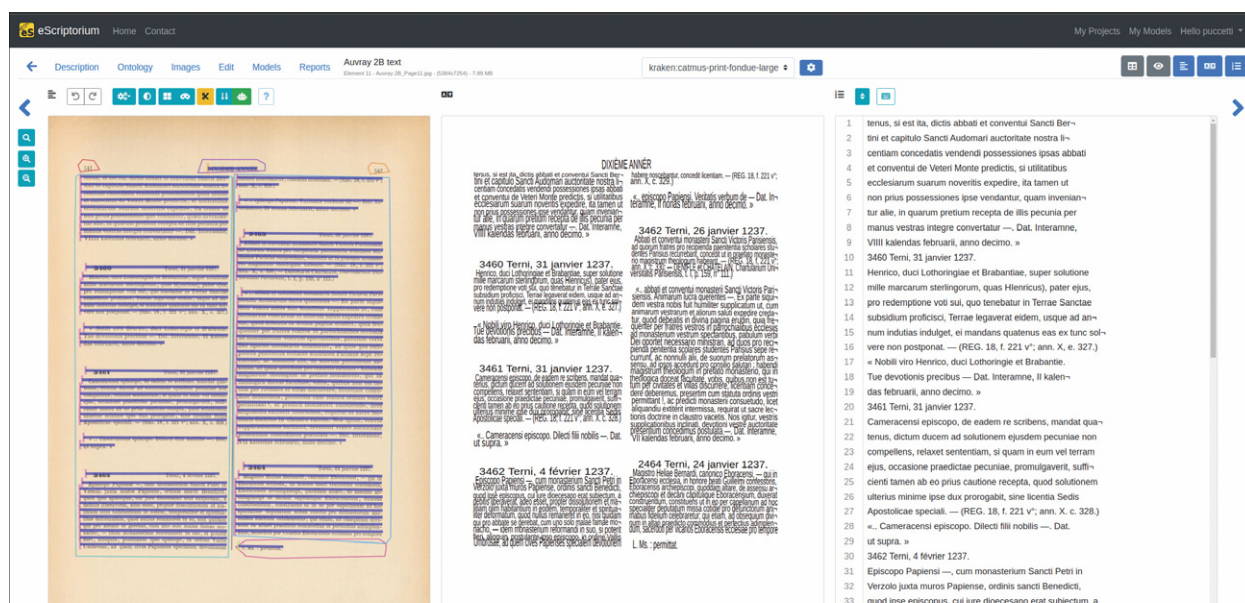


Figure 1: The annotation interface with the segmented regions and the output of the OCR process.




Figure 2: T-SNE plot showing the distribution of samples from the two collections MGH (dark) and Auvray (light).

2. Translation-Summarization (backtranslate): The model first translates the Latin full text into English, generates a summary in English, and then translates it back into Latin.

Quantitative evaluation using Rouge and Bleu metrics showed that GPT-4 outperformed the Llama models, achieving a Rouge-1 score of 0.39 on the MGH part of the dataset. Llama models performed better in the backtranslate setting, suggesting that the forced translation to English exploits the better summarisation capabilities of the model in their main language.

Qualitative analysis by a domain expert revealed challenges in accurately identifying key elements that are expected in regesta, such as the Pope's name, dates, and recipients. For example, GPT-4 correctly identified the Pope in 11 out of 20 cases and the recipient in 15, but only three dates were accurately extracted. These findings highlight the complexity of summarizing historical texts and the need for further refinement.

Conclusions

The REVERINO dataset provides a foundation for training models to automate the summarization of Latin texts, while the evaluation of LLMs offers insights into their strengths and limitations. Future work will expand the dataset to include more collections, further improve custom model training, and refine prompt engineering to improve summarization accuracy.

The ITSERR project is funded by the European Union - NextGenerationEU (CUP B53C22001770006).

Links:

[L1] <https://www.itserr.it/>

[L2] <https://zenodo.org/records/14971613>

References:

- [1] G. Puccetti et al., "REVERINO: REgesta generation VERSus latIN summarizatiOn", IRCDL 2025.
- [2] B. Kiessling, et al., "eScriptorium: an open source platform for historical document analysis", ICDARW, 2019.

Please contact:

Giovanni Puccetti, CNR-ISTI, Italy
giovanni.puccetti@isti.cnr.it

Who Are You, AI? – Artificial Intelligence as a Potential Metaphysical Humiliation of Humanity

by Michael Fleischhacker (University of Applied Sciences Burgenland), Igor Ivkić (University of Applied Sciences Burgenland and Lancaster University), and Friedrich Roithmayr (University of Applied Sciences Burgenland)

Artificial Intelligence as a potential fourth metaphysical humiliation—questioning human uniqueness in reason, freedom, and moral standing, and unsettling our place in the world after Copernicus, Darwin, and Freud.

Sigmund Freud identified three historical humiliations of humanity that have shaken the human sense of self [L1]. The first was the “cosmological humiliation” brought about by Nicolaus Copernicus' discovery that the Earth was not the centre of the universe, thereby depriving humans of their central position in the cosmos. The second humiliation, according to Freud, was the “biological humiliation” introduced by Charles Darwin's insight that humans were a product of evolution and not the creation of a higher being. Freud identified the third humiliation in the realm of the psyche and called it the “psychological humiliation”. In this view, the human subconscious undermines the supposed autonomy of the conscious mind: the “ego” is not, as Freud put it, “master in its own house”, but is driven by unconscious impulses. These three blows have effectively dethroned humanity from its self-appointed role as the centre and ruler of the world. In this article we explore whether AI is a potential fourth “metaphysical humiliation” of humanity, challenging the special status of humans as unique beings endowed with consciousness and reason, and the exclusivity of humanity as the only free and intelligent beings.

The emergence of Artificial Intelligence (AI) and the potential development of Artificial General Intelligence (AGI) poses a profound challenge to humanity's self-perception, touching on three critical dimensions. First, the domain of intelligence and problem-solving, long considered a hallmark of human uniqueness, is being redefined. AI systems such as AlphaGo and AlphaZero have demonstrated superhuman capabilities in complex tasks. AlphaZero, for example, mastered the strategy game Go by learning solely through playing itself, mimicking processes similar to human intuition and creativity. Such advances undermine the assumption that strategic thinking and intelligence are exclusively human traits. Second, it challenges the concept of consciousness and freedom. Jean-Paul Sartre famously described human beings as “condemned to be free”, with freedom arising from consciousness, which allows individuals to shape their own existence. If AI were to develop consciousness and subjective experience, it could fulfil Sartre's notion of freedom, thereby challenging humanity's monopoly on self-determination. This would force us to rethink our understanding of autonomy and concede that AI might also possess the capacity to “make itself”, as Sartre described it, dealing a profound blow to humanity's self-conception. Finally, the ethical and ontological supremacy tradition-




Figure 1: The Three Historical Humiliations of Humanity (Cosmological, Biological, Psychological) including AI as a potential Fourth Humiliation of Humanity (Metaphysical).

ally ascribed to human beings is at stake. Attributes such as dignity, morality, and responsibility have long been seen as uniquely human. However, the recognition of AI systems such as Sophia, who has been granted Saudi Arabian citizenship, as autonomous agents raises the question of whether AI will remain a mere tool or evolve into subjects with rights, responsibilities, and moral standing. Such developments force us to confront the unsettling possibility that humanity may no longer have exclusive claim to these defining qualities. Figure 1 shows the three historical humiliations of humanity according to Sigmund Freud, and adds AI as a potential fourth “metaphysical humiliation”.

Technological Developments as Catalysts for Metaphysical Humiliation

Advances in technology, particularly in the field of AI, are increasingly challenging humanity's metaphysical self-conception. One significant development is the rise of brain-computer interfaces (BCIs). These interfaces, such as implanted neural chips, allow direct communication between the human brain and AI systems. This unprecedented integration of human and machine intelligence is blurring the boundaries between the two and raising profound questions about what it means to be human. Similarly, the emergence of autonomous AI systems further complicates matters. These systems often operate through decision-making processes that are opaque to humans, a phenomenon known as the “black-box” problem. This lack of transparency undermines human control and accountability, leaving us to grapple with who or what is responsible for AI-driven decisions. Furthermore, transhumanist visions, such as achieving “digital immortality” by uploading consciousness or enhancing human cognition through AI, challenge the uniqueness of the biological body and the nature of human consciousness itself. Together, these developments are pushing humanity towards a metaphysical reckoning, forcing us to confront our diminishing exclusivity in intelligence, autonomy, and identity.

Counter-Arguments and Uncertainties

Despite these transformative possibilities, several counter-arguments highlight significant uncertainties in AI's potential to truly rival or surpass humanity. One is that AI still lacks subjective experience or consciousness in the human sense. Max Tegmark (philosopher) argues that while AI has made astonishing advances in processing power and functionality, replicating consciousness remains beyond current technological capabilities. In addition, ethical regulatory efforts provide a framework to mitigate the risks associated with the rapid development of AI. Initiatives such as the Asilomar AI Principles, formulated at the 2017 Beneficial

AI Conference by the Future of Life Institute, emphasise transparency, accountability, and the maintenance of human oversight over AI systems to prevent the disempowerment of humanity. Finally, sociologists point out that free will remains a uniquely human phenomenon, rooted in social and cultural contexts. While AI can simulate decision-making processes, it has yet to achieve the

emergent complexity that characterises human free will. These counter-arguments serve as a reminder that while AI challenges humanity's sense of self, its ultimate impact remains uncertain and dependent on both technological progress and ethical governance.

Conclusion – The Metaphysical Humiliation as a Cultural and Philosophical Challenge

The metaphysical humiliation of humanity by AI systems and AGI would not be a technical change, but an existential upheaval in the evolution of humanity. It would force humans to see themselves as part of a continuum with AI, rather than the crown of creation.

The emergence of AGI and superintelligent AI, endowed with consciousness, would logically imply that these AI systems would also be “condemned to be free” in Sartre's sense. In other words, such AI would also be forced to “create itself” rather than merely exist. Accordingly, Sartre's concept of freedom would still apply to humans, but it would also have to be extended to AGI and superintelligent AI. Whether this would ultimately constitute a fourth humiliation of humanity, following the three identified by Freud, remains a hypothetical question.

Ultimately, the responsibility lies with humanity: through human-centred AI development, ethical regulation, and critical reflection, this metaphysical humiliation can be mitigated or even prevented. Humanity now stands at a crossroads between self-assertion and the realisation that its supposed uniqueness in the world may not, in fact, be so unique after all.

Links:

- [L1] <https://doi.org/10.11588/diglit.25679#0012>
- [L2] <https://futureoflife.org/open-letter/ai-principles-german>

References:

- [1] M. Tegmark, “Leben 3.0: Mensch sein im Zeitalter Künstlicher Intelligenz”, H. Mania, Trans., 3rd ed., Ullstein, 2020.
- [2] J.-P. Sartre, “Das Sein und das Nichts: Versuch einer phänomenologischen Ontologie”, T. König and V. von Wroblewsky, Eds., 22nd ed., Rowohlt-Taschenbuch-Verlag, 2020.
- [3] S. Žižek, “Hegel im verdrahteten Gehirn”, 1st German ed., S. Fischer, 2020.

Please contact:

Michael Fleischhacker
University of Applied Sciences Burgenland, Austria
michael.fleischhacker@gmail.com

Green AI's Role in Cultural Heritage and Sustainability

by Enrico Barbierato and Alice Gatti (Catholic University of the Sacred Heart)

Cultural heritage represents societies' collective memory and identity, encompassing both tangible and intangible assets passed down through generations. As technology advances, Artificial Intelligence (AI) emerges as a powerful tool for preserving, analysing, and making heritage accessible to a wider audience. However, the environmental cost of AI, particularly the high energy demands of large-scale models, raises concerns about sustainability.

The project “linea di intervento D3.2, “Impatti scientifici ed etici delle applicazioni basate sull'Intelligenza Artificiale” (Scientific and Ethical Impacts of AI-Based Applications), developed internally across different departments within the Catholic University of the Sacred Heart in Brescia, was aimed at analysing the rapid development of Information Technology and its transformative impact on society, particularly in finance, healthcare, and industry. Initiated in 2019, the project spanned 24 months. It focussed on the role of AI, especially deep learning, in extracting knowledge from large datasets while addressing its interpretability and ethical implications. The project evaluated AI models not only for accuracy but also for their ethical considerations. One line of investigation focussed on the influence of large language models, such as ChatGPT, Google's Gemini, and Meta's LLaMA, which have demonstrated significant success. However, their carbon footprint is projected to be unsustainable in the long term.

The increasing energy consumption and environmental impact of AI models have led to a divide between “Red AI” and “Green AI,” a concept extensively analysed during the project. Specifically, Red AI refers to resource-intensive AI models that require enormous computational power, leading to high energy consumption and carbon emissions. In contrast, Green AI aims to optimise efficiency and reduce the environmental footprint by using smaller datasets, more efficient training techniques, or sustainable energy sources.

Figure 1 depicts the structure of a part of the project where we delved into a comprehensive, scientific survey [1] of the literature on Red and Green AI, categorising existing research into three clusters: the first study quantifies the environmental cost of Red AI and proposes Green AI as a solution; the second focuses on the carbon footprint of large deep learning models, and the third introduces methods to mitigate the computational cost of AI models. Our research highlights the trade-off between model performance and sustainability, questioning whether AI models can maintain high accuracy while reducing energy consumption. A major concern is that deep learning models, such as GPT and BERT, require

massive computational resources, often leading to diminishing returns—where increased complexity does not always translate into significantly better performance. Our study also explores the hardware and algorithmic factors that contribute to Red AI, such as computer architectures, inefficient data structures, and costly training methods. We review various techniques that improve efficiency, including pruning, quantization, and knowledge distillation, which help reduce model size and energy use while preserving accuracy. Additionally, the article presents case studies related to Natural Language Processing (NLP) models, such as BERT, GPT-3, and Google's Bard, emphasising their energy-intensive training requirements.

We claim that similar considerations apply to AI used for cultural heritage. AI models such as ChatGPT and generative AI can analyse, generate, and curate content related to historical artifacts, manuscripts, and artworks. One application is the digital restoration of ancient texts and paintings, where AI-powered image processing algorithms reconstruct damaged portions of manuscripts or art, restoring faded texts and pigments with remarkable accuracy. Moreover, AI-driven NLP models enable the translation and interpretation of ancient scripts, making historical documents accessible to a broader audience. Generative AI further extends these capabilities by reconstructing lost or fragmented cultural artifacts, generating historically accurate depictions of ancient cities, or simulating the voices of long-lost languages.

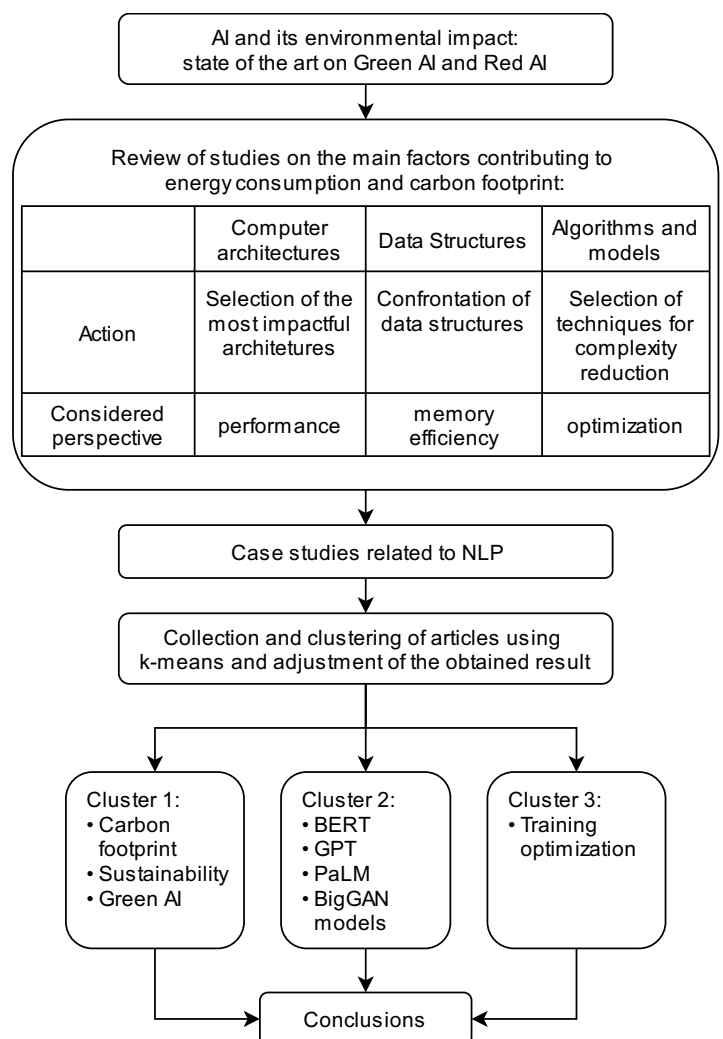


Figure 1: Motivation of the study.

Digital exhibitions allow global audiences to explore artifacts without physical constraints, fostering accessibility and inclusivity. The British Museum, for instance, offers a virtual tour of its extensive collection, leveraging AI to provide personalised recommendations based on user preferences. The Hereford Map [L1], an ancient cartographic masterpiece, serves as an example of how AI enhances historical cartography. By digitizing and analysing such maps, AI models can reconstruct historical geographical landscapes, provide interactive overlays comparing past and present terrains, and even predict lost historical sites based on textual references and spatial analysis. Another important application is the automation of metadata extraction in archival projects such as the Vatican's digitisation efforts, where deep learning techniques analyse vast collections of historical letters and religious manuscripts, enabling scholars to explore and contextualize these texts with unprecedented efficiency [L2].

The computational cost of training large AI models for cultural heritage applications is a critical consideration. The training of large language models such as ChatGPT-4 required approximately 25,000 Nvidia A100 GPUs for 90 to 100 days. This scale of operation translates to an estimated 50 gigawatt-hours (GWh) of energy consumption, underscoring the significant resources required for such advanced AI models. The environmental impact of such training procedures is substantial, as AI models necessitate high-powered computing clusters, often operating continuously for weeks. Compared to AI models in domains such as healthcare or finance, cultural heritage AI models might exhibit a unique computational complexity. Cultural heritage models engage in a mix of NLP, generative image synthesis, and 3D modelling, requiring a broader set of computations across multiple modalities. Consequently, the training complexity for AI applied to cultural heritage is likely to be higher than for traditional NLP tasks but lower than for high-dimensional reinforcement learning models.

Despite the energy demands, optimising AI for cultural heritage can follow principles of Green AI to ensure sustainable implementation. Techniques such as transfer learning, where pre-trained models are fine-tuned on smaller datasets, can significantly reduce computational requirements. Additionally, leveraging low-energy AI architectures, such as neuromorphic computing and quantised models, could further enhance efficiency. While AI models for cultural heritage do require extensive training, they can be optimised to minimise their carbon footprint while maximising their societal and educational impact.

Links:

[L1] <https://www.themappamundi.co.uk/>

[L2] <https://blogs.loc.gov/law/2020/10/collections-and-digitization-projects-of-the-vatican-apostolic-library/>

Reference:

[1] E. Barbierato and A. Gatti, "Toward Green AI: A methodological survey of the scientific literature," *IEEE Access*, vol. 12, pp. 23989-24013, Jan. 31, 2024.

Please contact:

Enrico Barbierato,
Catholic University of the Sacred Heart, Italy
enrico.barbierato@unicatt.it

Estimating Bias in Book Recommender Systems

by Savvina Daniil and Laura Hollink (CWI)

Online book recommender systems are known to propagate biases that historically exist in the book market, such as unfair treatment of authors based on their sensitive characteristics. At CWI, we study the challenges that researchers face in measuring and reporting on bias, by conducting extensive experimentation and exploring the gap between theory and practice.

Recommender systems is one of the AI applications that we interact the most with in our online journeys. From music, film and streaming, to news and social media posts, one thing is common: platforms attempt to improve user experience by automatically recommending content. Even public organizations such as libraries, where traditionally the content is curated, face challenges when users expect personalization in their online tools to support them when navigating the ever expanding information space.

Other than technical requirements that are always evolving, public organizations have to be mindful of potential negative consequences of training recommender systems on user-item interactions, such as bias. Bias refers to recommender models' tendency to find shortcuts and misinterpret the training data, and can lead to reinforcing stereotypes, malfunctioning for certain groups of people, recommending only items that are already well-known, and other effects. Bias is a valid concern in the context of library recommender systems, given known prejudices and unfair practices in the book market.

At the Human-Centered Data Analytics [L1] group of CWI and the Cultural AI Lab [L2], we wish to support KB, the National Library of the Netherlands, in their effort to adopt personalization while accounting for bias in their data. To this end, we evaluate bias propagated by frequently used recommender systems algorithms. We are specifically interested in popularity bias, the phenomenon where items that have been consumed by many users in the training data might be over-recommended and have their popularity further increase, regardless if the users actually liked them. Additionally, we pose that, given historically present biases, popularity is not random and might coincide with author sensitive characteristics such as gender, age, or nationality. In such cases, even if the system does not receive said characteristics as input, it will still over or under-recommend authors from a certain group as a direct result of popularity bias. For example, we found that in a benchmark dataset with book ratings American authors are over-represented, especially among the popular books. After preprocessing the dataset and training a set of collaborative filtering algorithms, we noticed that algorithms that propagate popularity bias also over-recommend American authors, as seen in Figure 1 [1].

In this process, we notice a set of challenges. For one, bias is a complex term, and while it is somewhat standardized when it comes to recommender systems, there are still critical decisions to make when measuring it. Given the inherently appli-




Figure 1: At CWI, we found American authors to be over-represented in a benchmark dataset with book ratings, especially among the popular books. After processing the dataset and training a set of collaborative filtering algorithms, some of them recommended up to 80% American authors.

cation-based character of recommendation, building on the rich output of library and information science can provide guidance in these decisions. Bridging theoretical and applied studies, while necessary, is ambitious. Differences such as divergent terminology obfuscate the process of accessing the right material from outside one's field: when computer scientists say 'bias', library and information scientists may say 'unjust treatment', 'marginalization', 'erasure' etc. For this reason, bias in book recommender systems research often lacks strong theoretical background, continuity, and even applicability in the real world. We are currently working on a survey paper in book recommendation that highlights this issue.

Another challenge with estimating bias is the dependence of the result on design and implementation details that often go unnoticed in research. The field of recommender systems is known to suffer from reproducibility issues, which also harm continuity. In our extensive reproduction work, we evaluated popularity bias propagation by a set of collaborative filtering methods such as Nearest Neighbor-based, Matrix Factorization, and Neural Network-based. We found that studies that measure popularity bias in different domains and attempt to reproduce each other adopt different evaluation strategies that ultimately lead to divergence in reported bias [2]. In addition to evaluation strategy, through data synthesis and simulation we found that implicit dataset properties interact with algorithm configurations that differ between widely used recommender systems frameworks and produce different conclusions regarding whether an algorithm or dataset is prone to exhibiting popularity bias [3]. We conclude that researchers have to be careful when simulating a system, as well as be specific and avoid generalized conclusions outside the boundaries of their studies.

Bias estimation is a valuable and necessary step in the development and adoption of complex learning-based systems. As a next step, we intend to move on from simulations and benchmark datasets to evaluating bias in the loaning patterns of users of the public libraries in the Netherlands and Denmark, in collaboration with DBC Digital. We hope that

the knowledge we've built will support the effort for effective communication and interdisciplinary dialogue between researchers and library practitioners to ensure that AI systems benefit society.

Links:

- [L1] <https://kwz.me/hxy>
[L2] <https://www.cultural-ai.nl/>

References:

- [1] S. Daniil, et al., "Hidden author bias in book recommendation," arXiv preprint arXiv:2209.00371, 2022.
- [2] S. Daniil, et al., "Reproducing popularity bias in recommendation: The effect of evaluation strategies," ACM Transactions on Recommender Systems, vol. 2, no. 1, pp. 1–39, 2024.
- [3] S. Daniil, et al., "On the challenges of studying bias in Recommender Systems: The effect of data characteristics and algorithm configuration," Information Retrieval Research, vol. 1, no. 1, pp. 3–27, 2025.

Please contact:

Savvina Daniil, CWI, The Netherlands
s.daniil@cwi.nl

Laura Hollink, CWI, The Netherlands
l.hollink@cwi.nl

Cultural Contexts in AI Design: Addressing Cultural Biases in AI Models and Data

by Javier Ibanez-Guzman, Imane Taourarti, Arunkumar Ramaswamy (Renault Group) and Ayesha Choudhary (Jawaharlal Nehru University)

This article explores how AI systems in autonomous vehicles perform differently across regions due to cultural, geographical, and socioeconomic factors embedded in their training data.

Modern vehicles increasingly rely on advanced driving assistance systems (ADAS) and autonomous functions, deploying AI-Enabled Systems (AIES) to enhance safety and passenger comfort. These systems have gained significant support worldwide, driven by legislation and evaluation programs like EuroNCAP, which promotes their adoption and enhance consumer confidence [L1]. These AIES systems, especially in the perception functions, rely heavily on data-driven methods. Machine learning (ML) models are trained using large datasets collected through extensive field campaigns, enabling detection, classification, and spatial understanding critical for decision-making. Rigorous testing follows to evaluate performance, focusing on metrics like false positives and false negatives for safety-critical functions such as pedestrian detection, lane departure warnings, and collision avoidance systems.

The behaviour of a vehicle reflects very much the context in which it evolves. The state of the road network, the complexity of the traffic sharing it (density, heterogeneity of vehicles, adherence to formal rules), the weather conditions affecting driving conditions (precipitation, visibility, temperature extremes), the quality of the infrastructure, etc. The different traffic

agents with which it interacts, the presence of animals or debris, etc. As shown in Figure 1, we can state the traffic agents reflect other conditions like social norms, law compliance, the manner they communicate, the driving skill level / training, etc. An important factor is the notion of risk, this reflect the economic needs of the different traffic agents, safety valuations, and cultural attitudes of the different traffic agents. This short analysis provides an insight into how locality reflects vehicle behaviour that is captured when logging data used to derive the machine learning models found in ADAS or AD applications [1]. This granularity provides what we term the locality to data. As a result if the AD systems are trained in a given locality, which is substantial different to the one that it is to be deployed, for the systems to respond in the same manner will be challenging unless mechanisms are provided to address the bias that is reflected in the data.

The main issue is that the ADAS and AD systems that incorporate functions that are derived from data their performance is significantly influenced by the datasets used for training [2]. These datasets often originate from specific regions, usually in northern hemisphere countries with distinct driving patterns, infrastructure, laws, and environmental conditions. In contrast, driving behaviours and road conditions in southern hemisphere countries, particularly in low- and middle-income regions, differ significantly due to ethnographic, legal, and geographical factors.

This geographical and cultural mismatch introduces biases in AI models, resulting in systems that may not perform effectively or safely when deployed in areas with differing local conditions. For example, a model trained on data from European roadways might struggle to adapt to the chaotic traffic patterns or poorly maintained infrastructure often found in other regions. The consequences of these biases can be severe, particularly in terms of safety outcomes, potentially reinforcing existing disparities in road safety between regions rather than mitigating them.

Addressing these biases is crucial to ensuring equitable and reliable road safety solutions worldwide. By incorporating di-

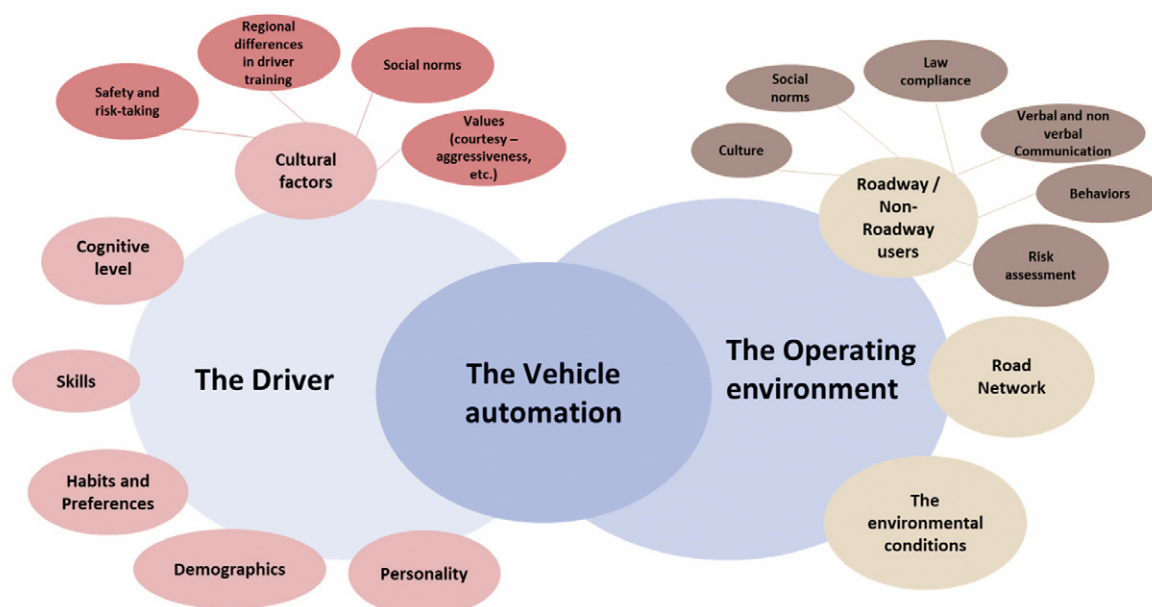


Figure 1: Interconnection between Vehicle, driver, and operating environment with their contributing factors.

verse datasets that reflect local driving conditions, cultural norms, and behavioural patterns, automotive AIES can be better equipped to handle the complexities of global deployment. The development of culturally adaptive autonomous systems represents not merely a technical challenge but an ethical imperative to ensure that the benefits of advanced vehicle safety technologies are equitably distributed across diverse global populations.

Links:

[L1] <https://kwz.me/hFv>

References:

- [1] I. Taourati et al., “Towards cross-cultural intelligent vehicles: A review,” in *Proceedings of VEHICULAR 2023, The 12th Int. Conf. on Advances in Vehicular Systems, Technologies and Applications*, 2023.
- [2] V. Diviš, T. Schuster, and M. Hruz, “Domain-centric ADAS datasets,” in *SafeAI 2023: The AAAI’s Workshop on Artificial Intelligence Safety*, 14 Jan. 2024.

Please contact:

Imane Taourati
Research Department/Renault Group, France
imane.taourati@ampere.cars

From Toxic Memes to Toxic Symbols and Back - Towards Automating Toxic Meme Analysis

by Davide Ceolin (CWI), Delfina Sol Martinez Pandiani (University of Amsterdam), and Erik Tjong Kim Sang (NL eScience Center)

Memos are extremely viral information items that, in many cases, are used to spread hateful messages. CWI, in collaboration with the University of Amsterdam and the Netherlands eScience Center investigates the effectiveness of semi-automated methods in analysing toxic memes.

This research is conducted by CWI, in collaboration with the University of Amsterdam and the Netherlands eScience Center, in the context of the Eye of the Beholder project, which is funded by the Netherlands eScience Center.

In the era of rising multimodal content consumption, the interplay between visual and textual elements, especially within computer science, is a key research focus. Notably, there's a growing interest in automatically discerning semiotic relationships between images and text. Simultaneously, a growing body of research delves into potentially problematic multimodal content in social media, particularly due to users' ease of engaging with and believing misinformation in multimodal formats. This underscores the heightened interest in automatically detecting hateful, propagandistic, fake, or otherwise mis-

leading or manipulative content in multimodal memes and social media posts.

To this aim, using the PRISMA methodology, we surveyed 158 papers and identified 30 datasets that focus on the problem of hateful meme classification [1]. From this analysis, we observe that in the literature, meme toxicity is characterized in rather diverse and misaligned manners. To this aim, we propose a harmonization of the terminology used, and we also introduce a framework to characterize such terminology. In particular, we identify three content-based dimensions of meme toxicity: target, intent, and conveyance tactics. The framework that we propose illustrates the relationships between these dimensions and meme toxicities. This allows characterizing the current classes of toxic content, as well as provides the means to represent new classes in case these emerge.

Having defined such a framework, we investigated the possibility of using this harmonized classification to automate the task of meme analysis. In particular, we identified in LLMs a potential tool for this goal, because these models can leverage large amounts of contextual information and can help discern the multiple layers of information that memes often portray. We performed preliminary analyses on using LLMs for this aim [2]. In particular, given the amount of resources required to run LLMs, we tested whether locally run LLMs can perform satisfactorily well in assessing the toxicity of memes. The results we obtain are promising: while it is true that larger LLMs perform better in assessing the toxicity of memes, it is also true that locally run models still show promising performance. In particular, locally run LLMs show promising results both in the assessment of the toxicity of memes and in the identification of toxic symbols that these contain.

Several factors can affect performance when analyzing toxic memes with LLMs. In particular, memes contain several layers of often implicit semantics, which may be hard to grasp automatically because of their complexity or because of the amount of specific background information needed to understand them. Moreover, guardrails introduced to avoid the generation of toxic content may impede the analysis of toxic items. The next work we developed is propaedeutic to addressing the former problem.

An important source of knowledge for toxic content detection is represented by The Global Extremist Symbols Database, which contains those symbols that were classified as toxic by The Global Project Against Hate and Extremism (GPAHE). In fact, while memes are complex objects composed of multiple semantic layers and contain implicit messages, many toxic memes portray their toxicity through the explicit use of toxic symbology. For this reason, in a subsequent work [3], we propose OnTox and OnToxKG, an ontology and a knowledge graph to characterize potentially toxic symbols from the Global Extremist Symbol Database. We demonstrate that the use of this symbolic representation of these symbols, often used in toxic memes, is beneficial for the task of toxicity detection. OnToxKG contains multimodal data of almost 800 potentially toxic symbols, and the use of Linked Data to represent them allows to easily link them to generic sources like Wikidata (<https://wikidata.org/>) and WordNet (<https://wordnet.princeton.edu/>), to facilitate disambiguation and contextualization of the symbols, as demonstrated also through a set of

SPARQL queries. Via such queries, we show that it is easy to use the ontology and the knowledge base to retrieve, for instance, symbols representing similar concepts, as well as retrieve memes containing specific symbolic representations (assuming that the meme itself was previously analysed to identify the relevant symbols).

In the future, we plan to further advance the analysis of memes in a (semi-)automated manner, exploring the possibility of automating this task by OnToxKG in combination with Linked Open Data sources and LLMs, considering also the potential subjectivity of this task.

References:

- [1] D. S. Martinez Pandiani, E. Tjong Kim Sang, and D. Ceolin. Toxic Memes: A Survey of Computational Perspec-

tives on the Detection and Explanation of Meme Toxicities, <https://arxiv.org/abs/2406.07353>, 2024. [Under review]

- [2] E. Tjong Kim Sang, D. S. Martinez Pandiani, and D. Ceolin. Evaluating Locally Run Large Language Models on Toxic Meme Analysis. 25th International Conference on Web Engineering (ICWE 2025) 2025 [to appear]
- [3] D. S. Martinez Pandiani, E. Tjong Kim Sang, and D. Ceolin. OnToxKG: An Ontology-Based Knowledge Graph of Toxic Symbols and their Manifestations. 25th International Conference on Web Engineering (ICWE 2025) 2025 [to appear]

Please contact:

Davide Ceolin, CWI, The Netherlands
Davide.Ceolin@cwi.nl



University
of Cyprus

Professor Position in AI XM Chair at the University of Cyprus

The Department of Computer Science at the University of Cyprus [L1] announces one academic position at the rank of Professor in the field of Artificial Intelligence. The position will be filled under XM's designated chair entitled "The XM Chair in Artificial Intelligence".

The Chair is funded by TP Servglobal Limited, a member of the XM Group [L2]. XM is a leading provider of online investment services internationally. Upon appointment, the Chair's holder will receive a €30,000 research grant from the XM Group, in addition to their initial funding under the University of Cyprus Regulations. This funding will support research development, promotional activities, and collaboration with the XM Group and other relevant international research centers and universities. The Chair's holder will play a key role in advancing teaching and research in Artificial Intelligence, benefiting the academic community, society, and the economy.

For this vacancy, a university degree and a PhD title from an accredited university are required. The minimum re-

quirements for the academic rank of a Professor are set at <https://rb.gy/k3pel>

In accordance with the applicable legislation, the annual gross salary (including the 13th salary) for full-time employment is: Professor (Scale A15-A16) €80.094,70 - €104.114,56.

Employee contributions to the various State funds will be deducted from the above amounts.

The selected candidate will be expected to teach in both Greek and English within the Department of Computer Science's undergraduate and graduate programs. Therefore, proficiency in both Greek and English is required. Candidates do not need to be citizens of the Republic of Cyprus.

Candidates are invited to submit their applications electronically by uploading the following documents in English and in PDF format at the following link: <https://applications.ucy.ac.cy/recruitment>

- 1 Cover letter
- 2 Curriculum vitae
- 3 Copy of ID/Passport

4 Copy of a bachelor's degree from an accredited University is required

5 Copy of a PhD degree from an accredited University is required

6 Review of previous research work and a brief description of future research projects (up to 3 pages)

7 List of publications

8 Representative publications (up to 3 publications which should be submitted separately).

9 The names and email addresses of three professors from whom confidential letters of recommendation (in English) will be automatically requested upon submission of the application. Letters of recommendation may be submitted up to 7 days following the deadline for submission of applications. It is the responsibility of each candidate to ensure that the references are submitted. In case the letters of recommendation are not submitted on time, the application will not be considered.

The deadline for applications is Friday the 4th of July 2025.

A more detailed announcement is available at:
<https://tinyurl.com/XMChairAI>

Development of Advanced Intelligent Robot Platform for Industrial Applications

by János Hollósi, Rudolf Krecht, and Áron Ballagi
(Széchenyi István University)

Deploying traffic cones has never been smarter or safer. This article unveils a cutting-edge autonomous robot platform equipped with advanced GNSS, AI-driven vision, and precise robotic manipulation to revolutionize traffic cone placement and retrieval. With a design tailored for efficiency, safety, and centimeter-level accuracy, this innovation tackles monotony and danger head-on, setting a new benchmark in automated traffic management.

Autonomous systems have consistently demonstrated their efficacy in tackling tasks that are inherently dirty, dull, and perilous across various sectors. Among these tasks, the deployment of temporary road signs and other traffic management devices distinctly embodies these three characteristics. Consider the repetitive chore of setting up traffic cones—often executed under challenging weather conditions and dangerously close to high-speed vehicular traffic. Moreover, certain specialized activities, such as setting up racetracks or vehicle testing tracks, demand placement of traffic cones with centimeter-level accuracy.

To address the need for both monotony and precision in traffic cone deployment and collection, the Vehicle Industry Research Center [L1] and the Autonomous and Intelligent Robotics Laboratory [L2] of Széchenyi István University introduce an innovative autonomous mobile robot platform designed specifically for this application. This robust platform combines a skid-steer robot base with an integrated robotic manipulator. Positioning is achieved through a Real-Time Kinematic Global Navigation Satellite System (GNSS-RTK). In case of collection, traffic cone detection utilizes a blend of camera systems and a bespoke artificial intelligence-driven machine vision solution. Additionally, the platform is equipped with a safety laser scanner for detecting obstacles.

Operation of this autonomous platform is facilitated through a tailored web-based user interface. Users can designate the precise locations for cone placement on a digital map. Subsequently, the robot systematically places the cones using GNSS coordinates. In a similar manner, the collection of traffic cones is also automated, utilizing both GNSS data and inputs from the machine vision system, thereby streamlining the process and enhancing safety and efficiency on the roads.

For the precise manipulation of traffic cones at locations specified by GNSS, the development of a robust, heavy-duty mobile robotic platform is indispensable (Figure 1). This platform must be equipped to transport approximately 30 traffic cones, collectively weighing around 150 kilograms, in addition to housing all the necessary apparatus for obstacle avoidance, as well as the automated placement and retrieval of traffic cones. Given the paramount importance of load capacity and stability, coupled with the requirements for controllable



Figure 1: Our robot platform under development.

low-speed maneuverability, the decision was made to engineer a four-wheel drive, skid-steer robot platform outfitted with air-inflated rubber tires. The platform is powered by onboard batteries [1].

The architectural framework of the platform's onboard systems is organized into six fundamental categories: power management, propulsion, manipulation, machine vision, low-level control, and high-level control. Further enhancing its functionality, the platform incorporates a modular low-level control system. This system receives input data from a sophisticated high-level computational system composed of two computers operating under ROS 2 (Robot Operating System 2). One of these computers is dedicated to task-specific functions related to machine vision, whereas the other serves as the central controller, orchestrating the integration of all systems within the robot platform. For traffic cone manipulation, a collaborative robotic arm equipped with a custom electric gripper is used. This collaborative robotic arm was chosen based on two considerations: the platform's potential future use for other tasks necessitates a multipurpose manipulator, and the need to operate safely near human workers, where collaborative functions provide the safest solution for mobile applications.




Figure 2: Visualization of the top of the detected cones in the space seen by the stereo camera.



Figure 3: The robot arm, equipped with a stereo camera, scans the robot platform's environment and positions itself to collect the cone.




Figure 4: The cones and their tops detected by our AI model.

A key element of our mobile robot platform is the artificial intelligence-based machine vision system. This module is designed to extract crucial information about the robot's surroundings through the analysis of data from various sensors, such as cameras and lidars (Figure 2). It does this using advanced artificial intelligence technologies, such as custom-developed neural networks, along with established image processing and machine vision methods. The objective of these intelligent software modules is to enable the detection of objects, segmentation of the environment, and estimation of distances and spatial positions of obstacles in front of the robot, based on sensor data. The purpose of this system is to provide valuable input to the other components of the robot. This includes, for instance, the robot's route planning algorithms, or the control software of the robot arm integrated on the platform.

Our current development is focused on enabling the robot to recognize roadside traffic cones for automatized collection (Figure 3). To this end, an artificial neural network has been developed which is optimized for our target hardware, which is a Connect Tech Rudi-AGX Embedded System with NVIDIA® Jetson AGX Xavier™. We have considerable experience in the development of systems of this nature, which has enabled us to create an effective solution to this problem.

During the cone collection process, the robot arm is positioned above the cones, allowing the camera integrated into the robot arm to provide a top view (Figure 4) of the cones. In order for the neural network to be able to recognize the cones from such a view, a custom dataset had to be developed, as no such solution is currently available on the market. Our custom-built dataset contains diverse top-view images of traffic cones in varying environmental conditions, where each image is individually labelled for training AI models. Our neural network

was trained using this unique dataset and is thus able to detect the cones with high efficiency [2].

The research was supported by the European Union within the framework of the National Laboratory for Artificial Intelligence. (RRF-2.3.1-21-2022-00004).

Links:

[L1] <https://jkk-web.sze.hu/?lang=en>

[L2] https://szolgaltatas.sze.hu/en_GB/autonomous-and-intelligent-robotics-laboratory

References:

[1] R. Krecht, Á. Ballagi, "Machine Design of a Robot Platform for Traffic Cone Handling" ["Tesztpálya kiszolgálására alkalmas autonóm robotplatform gépészeti tervezése"] in Conference Proceedings of Mobility and environment – Future-Shaping Automotive Research ["Mobilitás és környezet – Jövőformáló Járműipari Kutatások"], 2022.

[2] J. Hollósi, N. Markó, and Á. Ballagi, "Position and Orientation Determination of a Traffic Cone Based on Visual Perception" ["Forgalomtechnikai terület bójá vizuális észlelésen alapuló pozíció és orientáció meghatározása"] in Conference Proceedings of Digital Automotive Research at Széchenyi István University ["Digitális Járműipari Kutatások a Széchenyi István Egyetemen – Konferenciakiadvány"], 2021.

Please contact:

Áron Ballagi
Széchenyi István University, Hungary
ballagi@ga.sze.hu



SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Call for Proposals

Dagstuhl Seminars and Perspectives Workshops

Schloss Dagstuhl – Leibniz-Zentrum für Informatik is accepting proposals for scientific seminars/workshops in all areas of computer science, in particular also in connection with other fields.

If accepted, the event will be hosted in the seclusion of Dagstuhl's well known, own, dedicated facilities in Wadern on the western fringe of Germany. Moreover, the Dagstuhl office will assume most of the organisational/ administrative work, and the Dagstuhl scientific staff will support the organizers in preparing, running, and documenting the event. Thanks to subsidies the costs are very low for participants.

Dagstuhl events are typically proposed by a group of three to four outstanding researchers of different affiliations. This organizer team should represent a range of research communities and reflect Dagstuhl's international orientation. More information, in particular details about event form and setup, as well as the proposal form and the proposing process, can be found on

<https://www.dagstuhl.de/dsproposal>

Schloss Dagstuhl – Leibniz-Zentrum für Informatik is funded by the German federal and state government. It pursues a mission of furthering world class research in computer science by facilitating communication and interaction between researchers.

Important Dates

- *Next submission period:*
October 15 to November 1, 2025
- *Seminar dates:*
Between October 2026 and September 2027 (tentative).

The Importance of Software

A new position paper authored by an ERCIM task group highlights the crucial role of software technology, development, and maintenance in securing Europe's economic prosperity. The document calls for greater investment and attention to this often-overlooked area, stressing its foundational importance across industries and innovation.

Current Challenges

Europe's digital sovereignty strategy currently emphasizes hardware, neglecting vital investments in higher-level software. While European initiatives in chip design and hardware are essential, the most significant economic value increasingly originates from software-centric platforms. Europe's insufficient software capabilities make it dependent on foreign solutions. Without strengthening its own software ecosystem, Europe risks commoditisation of its hardware investments and loss of strategic autonomy.

Intermediate high-level software—including libraries, frameworks, and operating systems—is critical to developing competitive applications. However, industry typically regards this foundational software as non-value-adding, leading to underfunding despite its crucial role. Consequently, no major production-grade operating system is currently developed in Europe, creating a strategic gap in Europe's technology stack.

Experimental software research, essential for innovations, also suffers from inadequate support. Such research typically receives funding as a secondary element within larger projects, limiting researchers' capacity for genuinely explorative work. Europe's software R&D expenditure is considerably lower than competitors, approximately one-tenth of the US and one-third of China, threatening Europe's competitive position globally.

Vision for Europe's Software Future

The document proposes strategic shifts to secure Europe's digital sovereignty:

- **Software Research Leadership:**
Dedicated funding should enable high-risk, explorative software research, particularly leveraging AI advances (including generative AI).

Such focused research would foster transformative innovation, positioning Europe as a global leader in next-generation software development.

- **Software Production Leadership:**
Europe must invest in commodity software developed within the EU, ensuring alignment with European regulatory standards. Public funding for commodity software can effectively train software professionals, enhance industry innovation, and lower development costs, thus bolstering Europe's global competitiveness.
- **Software Maintenance Leadership:**
Sustained maintenance of critical software is vital, as software rapidly becomes obsolete without regular updates. Europe must actively maintain critical open-source software, ensuring longevity, continuous knowledge retention, and smooth technology transfer from research to industry.

Open-Source Software as a Strategic Asset

Open-source software significantly contributes to Europe's GDP, providing high returns by reducing vendor lock-in and facilitating collaboration between academia and industry. Nevertheless, Europe lags behind the US in open-source development, as US-based projects often receive substantial contributions from large corporations. Greater European involvement in global open-source governance would enhance strategic influence, benefiting European companies, startups, and SMEs.

Call to Action

To realize this vision, the document urges continuous and targeted public investment in software research, commodity software development, and maintenance. It specifically recommends matching new EU regulations on AI, sustainability, and privacy with dedicated software funding initiatives. Such sustained investment would help Europe reclaim a leadership role in software, promoting economic resilience and digital independence in a competitive global landscape.

The full paper, including the list of contributing task group members and references, is available from the ERCIM web site for download at <https://kwz.me/hGs>



Vanessa Evers Appointed New Director of CWI

The Board of NWO-I, the institute organization of NWO, appointed Prof. Vanessa Evers as Director of CWI, the national research institute for mathematics and computer science in the Netherlands. In September 2025, she will succeed the current director, Prof. Ton de Kok, who has led the institute since 2020 and will retire. Vanessa Evers is a leading researcher in the field of social artificial intelligence. The appointment as director of CWI is for a period of five years, with the possibility of a reappointment for the same period.

Vanessa Evers is a Full Professor and Chair of Human Media Interaction at the University of Twente's Computer Science department, a role she will continue alongside her position as CWI Director. She is also the founder and scientific co-director of the university's DesignLab. In addition, she serves on the Supervisory Board of Radboud University and is a visiting professor at Nanyang Technological University (NTU) in Singapore.

Beyond her scientific work, Vanessa Evers has been a strong advocate for diversity and inclusion in technology and academia. She actively promotes greater representation of women and underrepresented groups in STEM fields. Additionally, she emphasizes the ethical implications of technological advancements, advocating for socially responsible innovations that benefit all users.

Vanessa Evers will replace Ton de Kook as CWI's representative in ERCIM.

<https://www.cwi.nl/en/news/vanessa-evers-appointed-new-director-of-cwi/>



European Project Management

A European project can be a richly rewarding means of advancing your research or innovation activities to the state-of-the-art and beyond. Through ERCIM, our member institutes have participated in more than 100 European Union-funded projects in the ICT domain, conducting joint research activities while the ERCIM Office successfully manages the complexities of project administration, finances and outreach.

Horizon Europe: How can you get involved?

The ERCIM Office has recognised expertise in a full range of services, including:

- Identification of funding opportunities
- Recruitment of project partners (within ERCIM and through our networks)
- Proposal writing and project negotiation
- Contractual and consortium management
- Communications and systems support
- Organization of engaging events, from team meetings to large-scale workshops and conferences
- Support for the dissemination of results.

Please contact:
Peter Kunz, ERCIM Office
peter.kunz@ercim.eu



Endorsing Safeguarding, Protection & Provenance Management of Cultural Heritage

ENIGMA an EU-funded project, shall achieve excellence in the protection of Cultural Goods and artefacts from man-made threats by contributing to identification, traceability, and provenance research of Cultural Goods as well as by safeguarding and monitoring of endangered heritage sites.

Technologies of
Digital Identification &
Tracking Cultural Goods

Tools for
Decision Support &
Public Engagement



Funded by
the European Union

Start: 01 - 01 - 2023
End: 31 - 12 - 2025
Budget: 3 993 500 €
G.A.: 101094237



Project Coordinator:
The Aristotle University
of Thessaloniki

Consortium Partners



info@eu-enigma.eu



ERCIM – the European Research Consortium for Informatics and Mathematics is an organisation dedicated to the advancement of European research and development in information technology and applied mathematics. Its member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry.



ERCIM is the European Partner of the World Wide Web Consortium.



Consiglio Nazionale delle Ricerche
Area della Ricerca CNR di Pisa
Via G. Moruzzi 1, 56124 Pisa, Italy
www.iit.cnr.it



Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and Electrical Engineering, N 7491 Trondheim, Norway
<http://www.ntnu.no/>



Centrum Wiskunde & Informatica

Centrum Wiskunde & Informatica
Science Park 123,
NL-1098 XG Amsterdam, The Netherlands
www.cwi.nl



RISE SICS
Box 1263,
SE-164 29 Kista, Sweden
<http://www.sics.se/>



Fonds National de la
Recherche Luxembourg

Fonds National de la Recherche
6, rue Antoine de Saint-Exupéry, B.P. 1777
L-1017 Luxembourg-Kirchberg
www.fnr.lu



SBA Research gGmbH
Floragasse 7, 1040 Wien, Austria
www.sba-research.org/



Foundation for Research and Technology – Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
www.ics.forth.gr



Eötvös Loránd Research Network
Számítástechnikai és Automatizálási Kutató Intézet
P.O. Box 63, H-1518 Budapest, Hungary
www.sztaki.hu/



Fraunhofer ICT Group
Anna-Louisa-Karsch-Str. 2
10178 Berlin, Germany
www.iuk.fraunhofer.de



University of Cyprus
P.O. Box 20537
1678 Nicosia, Cyprus
www.cs.ucy.ac.cy/



INESC
c/o INESC Porto, Campus da FEUP,
Rua Dr. Roberto Frias, n° 378,
4200-465 Porto, Portugal
www.inesc.pt



UNIVERSIDAD DE MÁLAGA

Institute for Software Engineering and Software Technology
“Jose María Troya Linero”, University of Malaga
Calle Arquitecto Francisco Peñalosa, 18, 29010 Málaga
<https://gp.uma.es/itis>



Institut National de Recherche en Informatique
et en Automatique
B.P. 105, F-78153 Le Chesnay, France
www.inria.fr



I.S.I. – Industrial Systems Institute
Patras Science Park building
Platani, Patras, Greece, GR-26504
www.isi.gr



University of Warsaw
Faculty of Mathematics, Informatics and Mechanics
Banacha 2, 02-097 Warsaw, Poland
www.mimuw.edu.pl/