D4.2 Disseminating human-centric AI: overview of MUHAI dissemination activities

editors
Alessandra Fornetti
Folco Soffietti

MUHAI

Edited by Alessandra Fornetti and Folco Soffietti (Venice International University).
Published by Venice International University, Isola di San Servolo 30133 Venice, Italy

This project has received funding from the European Pathfinder Project under Grant Agreement N.951846.

D4.1 - Disseminating human-centric AI: overview of MUHAI dissemination activities

Contents

Foreword

The following chapters present the activities coordinated by the Venice International University (VIU) and executed by the MUHAI consortium. VIU built on years of established practice in science communication, developed through training programmes, collaborations, and projects—most notably QUEST and COALESCE. VIU ensured seamless implementation by subcontracting graphic design to a professional agency, social media management to a dedicated specialist, and web management to a webmaster. This structure enabled VIU to focus on coordinating these efforts, upholding the principles of science communication, and acting as a knowledge broker between AI researchers and professionals who could adapt complex content for wider audiences.

Throughout the project, MUHAI took the opportunity to study the public debates on AI and explore how researchers could contribute meaningfully to them. The insights gained from this work are outlined in the introduction and are detailed in Vol.3 of the MUHAI book.
AI communication unfolds within two distinct narratives: one shaped by experts and another driven by media enthusiasm. The rise of generative AI tools, such as ChatGPT, has heightened public interest, evoking both excitement and concern.

Tech enthusiasts and businesses often present an overly optimistic vision of AI's capabilities, portraying it as a revolutionary force. Conversely, artists, creative professionals, and younger generations frequently express scepticism, fearing job displacement and ethical risks. Meanwhile, the scientific community struggles to convey nuance, as corporate narratives and alarmist media reports often overshadow their voices.

A key challenge lies in bridging the gap between exaggerated promises and legitimate concerns, ensuring that public discourse remains evidence-based and balanced. MUHAI recognised that AI researchers have a crucial role in demystifying AI, highlighting its limitations and ethical considerations while countering misinformation. The project identified several key strategies for building trustworthy AI communication:

- Transparency: Clearly explaining how AI works, its strengths and limitations, and its implications for human labour.
- Ethical responsibility: Addressing concerns such as bias, privacy, and misinformation while actively debunking myths.
- Balanced messaging: Avoiding both hype and fearmongering, ensuring that policymakers and the public have realistic expectations.
- Collaborative storytelling: Engaging with communicators and journalists to create accessible narratives about AI research.

Despite these efforts, researchers alone cannot be responsible for public communication. Given their demanding workloads, the study suggests that integrating communication professionals within research teams would enhance outreach and ensure effective knowledge transfer. Effective AI communication requires a blend of traditional and innovative methods to engage diverse audiences. MUHAI identified several successful approaches:

- Humanising AI: Presenting AI as an assistive tool rather than a replacement, making it more relatable.
- Social media and video content: Short, focused videos and interactive formats proved particularly effective in reaching younger audiences.
- Public engagement through the arts: Collaborations with artists, exhibitions, and storytelling broadened societal discussions on AI's impact.
- University-led initiatives: Educational programmes and ethics courses reinforced responsible AI use among students.

MUHAI's Impact on Researchers' Awareness
In a series of interviews, the impact of MUHAI's communication work on AI researchers was explored. It emerged that for early-career researchers, MUHAI provided an eye-opening experience, demonstrating the significance of communication in shaping public perceptions of AI. Senior researchers, already familiar with these dynamics, reported only minor shifts in their perspectives.

Key takeaways from the project included:

- AI as a complement, not a replacement: MUHAI reinforced the need to shift AI narratives away from fear and towards coexistence with human intelligence.
- Rising public interest in AI: The emergence of Large Language Models (LLMs) increased demand for expert insights, encouraging researchers to engage more actively in public discussions.
- Gaps in EU communication metrics: Researchers expressed frustration that EU project metrics often fail to capture the true impact of communication efforts.

While social media engagement remained a time-consuming challenge, collaborations with artists and public exhibitions emerged as alternative and effective methods for AI outreach.

The Role of Communicators in Supporting AI Research
MUHAI underscored the importance of professional communicators in translating complex AI concepts into accessible narratives. Researchers highlighted the value of communicators in:

- Framing discussions and guiding outreach efforts.
- Refining scientific language to suit different audiences.
- Developing engaging content, including videos, blogs, and social media campaigns.

Corporate environments, such as SONY (project partner), demonstrated the advantages of integrating communicators directly within research teams. This approach allowed researchers to focus on their work while ensuring effective dissemination.

However, the project also identified a knowledge gap, and hence a challenge to overcome: many communicators lacked expertise in AI's technical aspects. Close collaboration with researchers is therefore essential to avoid oversimplification and ensure accurate representation. The following chapters will examine MUHAI's Dissemination and Ethics activities (T4.2 and 4.3) in relation to the project's initial objectives. This deliverable is the twin of one dedicated to Communication activities (T4.1).

Chapter 1: Building a community; Event participation

Events played a pivotal role in MUHAI's dissemination efforts, serving as key opportunities to engage the scientific community, liaise with policymakers, and explore applications in the food sector. The consortium actively sought out opportunities to present research findings, particularly at scientific conferences. MUHAI partners participated extensively in both European and global events, either by presenting research advancements developed within the project framework or by organising dedicated sessions and workshops.

Scientific Conferences and Workshops

MUHAI researchers contributed to numerous high-profile scientific gatherings, including but not limited to:

- The First International Conference on Hybrid Human-Artificial Intelligence (2021)
- IJCAI-ECAI 2022 | Workshop on Semantic Techniques for Narrative-Based Understanding
- CHI Play 2022
- The 21st IFIP International Conference on Entertainment Computing
- UM6P Science Week (2023)
- STAIRS - European Starting AI Researchers' Symposium (2023)
- CLIN 33 - Computational Linguistics in The Netherlands (2023)
- DRHA23 - Performing Cultural Heritage in the Digital Present
- The 61st Annual Meeting of the Association for Computational Linguistics
- BNAIC (2023 & 2024), ICAART (2024), ESWC (2024), ECAI (2024), AISCICOMM24, VECOMP (2024)

As a culmination of MUHAI's engagement with the scienfic community, the project's final scientific event was hosted at The 24th International Conference on Knowledge Engineering and Knowledge Management (EKAW-24), a reference event in the field. A dedicated tutorial was organised to introduce participants to MUHAI's human-centric AI approach. Attendees explored methodologies for building AI agents capable of understanding real-world activities through a hybrid approach, combining MUHAI outputs about:

Symbolic AI, including ontological modelling of narratives, automated knowledge extraction from large-scale knowledge graphs, and the storage of episodic and semantic memories.

Subsymbolic AI, encompassing language understanding, speech processing, reinforcement learning, and active learning.

Educational and Training Initiatives

Beyond scientific conferences, MUHAI actively contributed to education and capacity-building at an international level. Notable initiatives included the lecture "AI - Hype, Hysteria, and Reality: Towards Human-Centric and Empowering Digital Media", delivered at Chiang Mai University's College of Arts, Media, and Technology.

Additionally, the project co-organised several summer schools, including:

Interdisciplinary College (IK) 2021 – "Connected in Cyberspace"
International Semantic Web Research Summer School (ISWS 2023)

Summer school on Sentience and Responsibility in Critical Times, University of Pisa, Lama Tzong Khapa Institute in Pomaia, Italy. (2023)
International Semantic Web Summer School (ISWS 2023), held in Bertinoro, Italy.
Interdisciplinary College at Lake Möhne (Günne, 2024).

Industry Engagement and Commercialisation Opportunities

MUHAI actively pursued opportunities to disseminate its technological innovations to industry stakeholders, with a particular focus on the food sector. Business partners, such as APICBASE and SONY Computer Science Labs, participated in leading trade fairs and industry events, including:

- Casual Dining Show (London, 2021, 2022, 2023)
- Food Hotel Tech (Paris, 2021)
- Horeca Expo Gent (Ghent, 2021)
- GastroNord (Stockholm, 2021)
- Restaurant & Bar Tech Live (London, 2022)
- Horecava (Amsterdam, 2023 & 2024)
- FFCR (Stockholm, 2023 & 2024)
- Internorga (Hamburg, 2024)

MUHAI was also selected for participation in the EU Innovation Radar, an initiative that identifies high-potential innovations within EU-funded projects and supports their transition from research to market. As part of this initiative, MUHAI provided structured insights into the innovations developed within the project.

Engagement with Policymakers

MUHAI leveraged key events to engage policymakers and contribute to discussions on the future of AI in Europe. Examples include:

- AI4Belgium - Inspirational Session (2023), held at the European Parliament, exploring opportunities and risks for AI development in Belgium and Europe.
- MUHAI presentation at Bremen's State Media Authority (2023).
- Meeting with the Mayor of the State of Bremen (Online, 2023).
- 'AI for Social Change' Academy meeting in Brussels.
- Hearing with Helga Nowotny (2023), providing a statement for policymakers on AI governance.

Policymakers were also present at MUHAI's Public Final Event, held as part of i2b Bremen, which was attended by Bremen's Chancellor and Bremen's Senator for Environment, Climate, and Science.

Public Engagement and Outreach

Beyond academic, industry, and policy circles, MUHAI also engaged with the wider public through talks and cultural events. Notable examples include:

- The Milanesiana 2023.
- "Chat GPT: What Now?" – Public talk at ISI Foundation (Turin, Italy, 2024).
- Lecture at Letterenhuis (Antwerp, 2023).

04 December 2024

ECAI 2024

19th October 2024

04 December 2024

BNAIC 2024

18th November 2024

04 December 2024

What should AI do next? Symposium

13th November 2024

04 December 2024

The Limitations of Generative Artificial Intelligence

8th November 2024

03 December 2024

A Paradigm Shift in Computer Science?

28th November 2024

29 November 2024

MUHAI Newsletter #12 November 2024

30th November 2024

25 November 2024

REACT Workshop Explores Responsible and Ethical AI in CUI Technologies

18th November 2024

16 October 2024

VII Frontiers Scientific Dialogue

13th November 2024

13 September 2024

#11 MUHAI Newsletter released

July 2024

13 September 2024

Muhai final project meeting

12th-14th September 2024

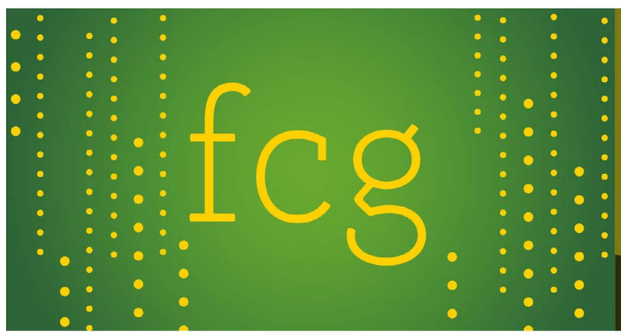

01 June 2021

Bringing Computational Construction Grammar into your Classes and Research

August 17, 2021 | Tutorial at the 11th International Conference on Construction Grammar

20 May 2021

Archives and AI: Coping with Climate Change

Workshop | May 22-24, 2021 @Venice International University | Chairman: Luc Steels. The workshop uses the AquaGranda Digital Community Memory as a source of...

14 April 2022

CHI Play 2022

November 2-5, 2022 | Bremen, Germany, and anywhere on earth

16 March 2022

IJCAI-ECAI 2022 | Workshop on semantic techniques for narrative-based understanding

July 24, 2022 | Organizers: Lise Stork (VUA), Katrien Beuls (Université de Namur), Luc Steels (VIU/BSC)

11 March 2022

MUHAI in the City of Light

March 14-16, 2022 | Organizers: Inès Blin and Remi van Trijp (Sony Computer Science Laboratories Paris)

29 November 2021

The first International Conference on Hybrid Human-Artificial Intelligence

June 13-17, 2022 | Deadline for abstract submission: February 25, 2022

13 August 2024

AI - Hype, Hysteria, and Reality

7th August 2024

06 August 2024

VECOMP 2024 - VALE Track

19th-24th October 2024

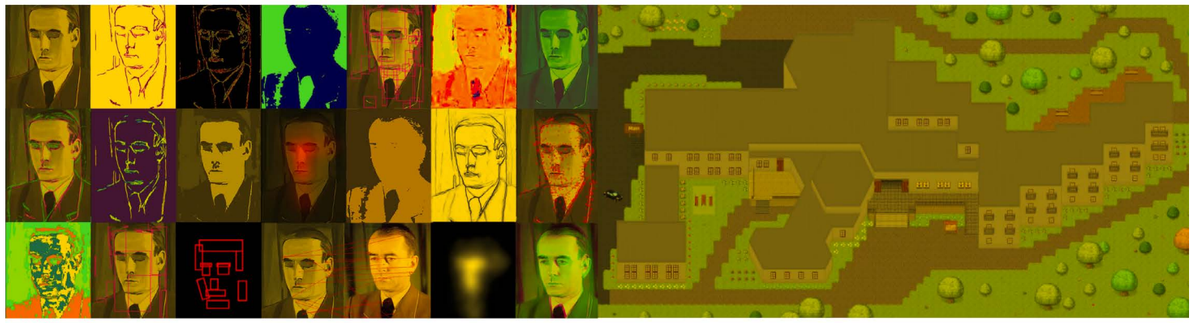

24 March 2021

SECRETS | Artificial Intelligence and Luc Tuymans

April 3 — May 2, 2021 Exhibition @ BOZAR Brussels What is the relationship between art and artificial intelligence (AI)? Can AI become a tool for the con...

12 March 2021

Interdisciplinary College | Annual Spring School

March 12 - April 9, 2021 | 'Connected in Cyberspace' is the 2021 edition of the Interdisciplinary College (IK)

10 March 2021

AI4BELGIUM | Citizen debate

March 18, 2021 | Key notes by Belgian PM De Croo, EU Director for Digital Industry & Artificial Intelligence Lucilla Sioli, and Luc Steels. 9.00-10.30 am - f...

16 January 2021

Taking Responsibility for Responsible Artificial Intelligence

December 16, 2020 | webinar Luc Steels, MUHAI Scientific Coordinator, talks of "History Lessons for the Future: Toward a Human-centric governance of AI". ...

Chapter 2: Building a Scientific Community; Papers

MUHAI's impact on the scientific community is reflected in the significant number -43- of academic papers produced throughout the project. These papers, authored either collaboratively between partners' teams or within individual teams, explore key topics directly related to MUHAI's research outputs and demonstrators.
To enhance accessibility and visibility, a dedicated section on the MUHAI website hosted the publications, with most papers also being promoted via social media upon release. The research was disseminated across various formats, including conference proceedings and journal articles, ensuring a broad academic reach.
A substantial portion of the research conducted within MUHAI contributed to the MUHAI Volumes, a comprehensive three-volume body of work that retraces the project's key findings and outputs. These volumes, curated by senior researchers Luc Steels, Robert Porzel, and Frank van Harmelen, were designed with a coherent visual identity, developed in collaboration with studio +fortuna. This structured publication approach ensured that MUHAI's research legacy would remain well-documented and accessible to the wider academic community. The volumes are available on Zenodo.

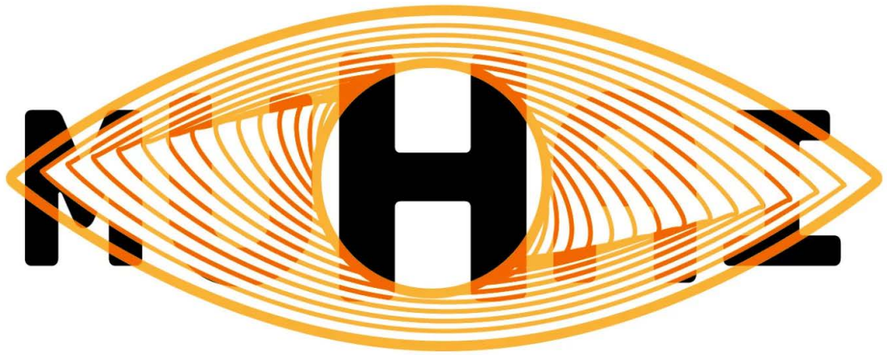

papers

Luc Steels. Norms and AI.. *N. Osman and L. Steels (Eds.): VALE 2023, LNCS 14520*, pages 1-7, 2024. PDF BibTeX

Katrien Beuls and Paul Van Eecke. Humans Learn Language from Situated Communicative Interactions. What about Machines?. *Computational Linguistics*, pages 1-34, August 2024. URL, DOI BibTeX

Jonas Doumen, Katrien Beuls and Paul Van Eecke. Modelling constructivist language acquisition through syntactico-semantic pattern finding. *Royal Society Open Science* 11:231998, 2024. URL BibTeX

Shiyao Zhang, Omar Faruk, Robert Porzel, Dennis Küster, Tanja Schultz and Hui Liu. Examining the Effects of Human-Likeness of Avatars on Emotion Perception and Emotion Elicitation. In *6th International Conference on Activity and Behavior Computing (ABC2024)*. 2024. URL BibTeX

Thomas Mildner Kamyar Javanmardi Robert Porzel Nima Zargham Vino Avanesi and Rainer Malaka. HASI: A Model for Human-Agent Speech Interaction. In *In Proceedings of CUI '24: ACM Conversational User Interfaces*. July 2024. URL BibTeX

Inès Blin, Lise Stork, Laura Spillner and Carlo Santagiustina. OKG: A Knowledge Graph for Fine-grained Understanding of Social Media Discourse on Inequality. In *Proceedings of the 12th Knowledge Capture Conference 2023*. 2023, 166–174. URL, DOI BibTeX

Jens Nevens, Robin de Haes, Rachel Ringe, Mihai Pomarlan, Robert Porzel, Katrien Beuls and Paul van Eecke. A Benchmark for Recipe Understanding in Artificial Agents. In Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti and Nianwen Xue (eds.). *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. 2024, 22–42. URL BibTeX

Lise Stork and Richard Zijdeman. MIRA-KG: A knowledge graph of hypotheses and findings for social demography research. 2023. URL BibTeX

Liesbet De Vos, Jens Nevens, Paul Van Eecke and Katrien Beuls. Construction grammar and procedural semantics for human-interpretable grounded language processing. *Linguist. Vanguard*, 0. URL BibTeX

Sofia Baroncini, Luc Steels and Remi van Trijp. Semantic data retrieval and integration for supporting artworks interpretation through integrative narrative networks(Short paper). In Antonis Bikakis, Roberta Ferrario, Stéphane Jean, Béatrice Markhoff, Alessandro Mosca and Marianna Nicolosi Asmundo (eds.). *Proceedings of the International Workshop on Semantic Web and Ontology Design for Cultural Heritage* 3540. 2023. URL BibTeX

Daniel Beßler, Robert Porzel, Mihai Pomarlan and Michael Beetz. Foundational Models for Manipulation Activity Parsing. In Timothy Jung, Claudia M tom Dieck and Sandra Maria Correia Loureiro (eds.). *Extended Reality and Metaverse*. 2023, 115–121. BibTeX

Mona Abdel-Keream, Daniel Beßler, Ayden Janssen, Sascha Jongebloed, Robin Nolte, Mihai Pomarlan and Robert Porzel. An Ontological Model of User Preferences. BibTeX

Robin Nolte, Mihai Pomarlan, Daniel Beßler, Robert Porzel, Rainer Malaka and John Bateman. Towards an Ontology for Robot Introspection and Metacognition. BibTeX

Laura Spillner, Nima Zargham, Mihai Pomarlan, Robert Porzel and Rainer Malaka. Finding Uncommon Ground: A Human-Centered Model for Extrospective Explanations. In *Proceedings of the 2023 IJCAI workshop on Explainable Artificial Intelligence (XAI)*. 2023. URL BibTeX

Rachel Ringe and Robert Porzel. Towards a Task-based Metric for Measuring Trust in Autonomous Robots for Everyday Activities. In *Proceedings of the CHI TRAIT Workshop on Trust and Reliance in AI-Assisted Tasks*. 2023. URL BibTeX

Jonas Doumen, Katrien Beuls and Paul Van Eecke. Modelling Language Acquisition through Syntactico-Semantic Pattern Finding. In *Findings of the Association for Computational Linguistics: EACL 2023*. 2023, 1317–1327. URL BibTeX

Paul Van Eecke, Jens Nevens and Katrien Beuls. Neural heuristics for scaling constructional language processing. *Journal of Language Modelling* 10(2):287–314, 2022. URL BibTeX

Mihai Pomarlan, Maria M Hedblom and Robert Porzel. Curiously exploring affordance spaces of a pouring task. *Expert Systems*, pages e13213, 2022. URL, DOI BibTeX

Laura Spillner John Bateman Thomas Mildner Robert Porzel Mihai Pomarlan and Carlo Santagiustina. Narrativizing Knowledge Graphs. In *Joint Proceedings of the 3th International Workshop on Artificial Intelligence Technologies for Legal Documents (AI4LEGAL 2022) and the 1st International Workshop on Knowledge Graph Summarization (KGSum 2022) co-located with the 21st International Semantic Web Conference (ISWC 2022). Virtual Event, Hangzhou, China, October 23-24*. 2022, 100–111. PDF BibTeX

Jens Nevens, Jonas Doumen, Paul Van Eecke and Katrien Beuls. Language Acquisition through Intention Reading and Pattern Finding. In *Proceedings of the 29th International Conference on Computational Linguistics*. 2022, 15–25. URL BibTeX

Inès Blin. Building a French Revolution Narrative from Wikidata. In *Proceedings of the IJCAI/ECAI Workshop on Semantic Techniques for Narrative-based Understanding*. July 2022. PDF BibTeX

Inès Blin. Building Narrative Structures from Knowledge Graphs. In *European Semantic Web Conference*. 2022, 234–251. URL PDF BibTeX

Luc Steels (ed.). *Foundations for Meaning and Understanding in Human-centric AI*. Venice International University, 2022. URL, DOI BibTeX

Luc Steels. Personal dynamic memories are necessary to deal with meaning and understanding in human-centric AI. In Alessandro Saffiotti, Luciano Serafini and Paul Lukowicz (eds.). *Proceedings of the First International Workshop on New Foundations for Human-Centered AI (NeHuAI) co-located with 24th European Conference on Artificial Intelligence (ECAI 2020), Santiago de Compostella, Spain, September 4, 2020* 2659. 2020, 11–16. URL BibTeX

Sinem Aslan and Luc Steels. Identifying Centres of Interest in Paintings Using Alignment and Edge Detection. In Alberto Del Bimbo, Rita Cucchiara, Stan Sclaroff, Giovanni Maria Farinella, Tao Mei, Marco Bertini, Hugo Jair Escalante and Roberto Vezzani (eds.). *Pattern Recognition. ICPR International Workshops and Challenges*. 2021, 589–603. URL BibTeX

Laura Spillner, Robert Porzel, Robin Nolte and Rainer Malaka. Conceptual Shadows: Visualizing Concept-specific Dimensions of Meaning in Word Embeddings with Self Organizing Maps. In *CEUR WORKSHOP PROCEEDINGS*. 2023. PDF BibTeX

Laura Spillner, Rachel Ringe, Robert Porzel and Rainer Malaka. "My, my, how can i resist you? " - examining user reactions to bogus explanations of ai. In Marianna Bergamaschi Ganapini, Andrea Loreggia, Nicholas Mattei, Francesca Rossi, Biplav Srivastava and Brent Venable (eds.). *Proceedings of the Workshop on Ethics and Trust in Human-AI Collaboration: Socio-Technical Approaches (ETHAICS 2023)* 3547. 2023. URL BibTeX

Lara Verheyen, Jérôme Botoko Ekila, Jens Nevens, Paul Van Eecke and Katrien Beuls. Neuro-symbolic procedural semantics for reasoning-intensive visual dialogue tasks. In *Frontiers in Artificial Intelligence and Applications*. Frontiers in artificial intelligence and applications series, IOS Press, 2023. BibTeX

Paul Van Eecke, Lara Verheyen, Tom Willaert and Katrien Beuls. The Candide model: How narratives emerge where observations meet beliefs. In *Proceedings of the The 5th Workshop on Narrative Understanding*. 2023, 48–57. URL BibTeX

Paul Van Eecke, Katrien Beuls, Jérôme Botoko Ekila and Roxana R\u adulescu. Language games meet multi-agent reinforcement learning: A case study for the naming game. *J. Lang. Evol.* 7(2):213–223, 2022. BibTeX

Katrien Beuls and Paul Van Eecke. Fluid Construction Grammar: State of the Art and Future Outlook. In *Proceedings of the First International Workshop on Construction Grammars and NLP (CxGs+NLP, GURT/SyntaxFest 2023)*. 2023, 41–50. URL BibTeX

L Verheyen, J Botoko Ekila, J Nevens, P Van Eecke and K Beuls. Hybrid Procedural Semantics for Visual Dialogue: An Interactive Web Demonstration. In *Poster session presented at Workshop on semantic techniques for narrative-based understanding - IJCAI-ECAI 2022*. 2022. BibTeX

Rachel Ringe and Robert Porzel. Towards a Task-based Metric for Measuring Trust in Autonomous Robots for Everyday Activities. In *Proceedings of the CHI TRAIT Workshop on Trust and Reliance in AI-Assisted Tasks*. 2023. URL BibTeX

Mihai Pomarlan and Robert Porzel. Narrative Objects. In *Proceedings of the IJCAI/ECAI Workshop on Semantic Techniques for Narrative-based Understanding*. July 2022. PDF BibTeX

Thomas Mildner Robert Porzel Laura Spillner Carlo R. M. A. Santagiustina. Towards Conflictual Narrative Mechanics. In *Proceedings of the IJCAI/ECAI Workshop on Semantic Techniques for Narrative-based Understanding*. July 2022. PDF BibTeX

Remi van Trijp, Katrien Beuls and Paul Van Eecke. The FCG Editor: An innovative environment for engineering computational construction grammars. *PLOS ONE* 17(6):1-27, 2022. URL PDF, DOI BibTeX

Maria Hedblom Mihai Pomarlan Vanja Sophie Cangalovic Johannes Pfau John Bateman M A Sebastian Höffner Robert Porzel and Rainer Malaka. Deep Understanding of Everyday Activity Commands for Household Robots. *Semantic Web Journal: Special Issue on Semantic Web for Industrial Engineering: Research and Applications*, 2022. PDF BibTeX

Nikolaos Kondylidis, Jie Zou and Evangelos Kanoulas. Category Aware Explainable Conversational Recommendation. , 2021. URL PDF, DOI BibTeX

Laura Spillner and Nina Wenig. *Talk to Me on My Level – Linguistic Alignment for Chatbots*. Association for Computing Machinery, 2021. URL PDF BibTeX

Maria Hedblom M Mihai Pomarlan and Robert Porzel. Panta Rhei: Curiosity-Driven Exploration to Learn the Image-Schematic Affordances of Pouring Liquids. In *Proceedings of the 29th Irish Conference on Artificial Intelligence and Cognitive Science*. 2021. PDF BibTeX

Robert Porzel. On Formalizing Narratives. In *Proceedings of the JOWO – Ontology Workshops*. 2021. PDF BibTeX

Robert Porzel Rainer Malaka Michael Beetz Maria Hedblom Mihai Pomarlan. Dynamic Action Selection Using Image Schema-Based Reasoning for Robots. In *Proceedings of the JOWO – Ontology Workshops*. 2021. PDF BibTeX

Vanja Sophie Cangalovic, Robert Porzel and John A Bateman. Streamlining Formal Construction Grammar. In *Proceedings of the 11th International Conference on Construction Grammar*. 2021. PDF BibTeX

Vanja Sophie Cangalovic Robert Porzel and John A Bateman. Filling Constructions: Applying Construction Grammar in the Kitchen. In *Proceedings of the 11th International Conference on Construction Grammar*. 2021. PDF BibTeX

Chapter 3: Maintaining the community; Newsletter

The MUHAI quaterly newsletter served as a concise yet effective way to keep stakeholders informed about the project's progress. With around 110 subscribers, the newsletter provided a digest of key activities and results every four months.
Each edition followed a structured format, ensuring consistency and engagement:
A seasonal banner, uniquely designed for each issue.
The "Word of the Month", highlighting a key concept related to MUHAI's research.
Dedicated sections showcasing new research outputs and papers.
Recaps of past events and announcements for upcoming conferences, workshops, and initiatives.
This structured approach helped maintain a clear and engaging communication channel with the MUHAI community, offering insights into both the academic and practical advancements made within the project.

Newsletter Archive

MUHAI Newsletter #13
March 2025

- March #wordofthemonth
- Outputs > Volume 3
- Event > EKAW 2024
- Event > Women in science
- New scientific papers
- New tutorials > #TrAIning

MUHAI NL #12
MUHAI NL #11
MUHAI NL #10
MUHAI NL #9
MUHAI NL #8
MUHAI NL #7
MUHAI NL #6
MUHAI NL #5
MUHAI NL #4
MUHAI NL #3
MUHAI NL #2
MUHAI NL #1

Click here to view it online

MUHAI Newsletter #11 | July 2024

July #wordofthemonth

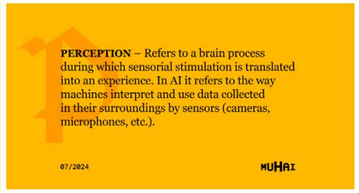

Check out the previous words of the month!

Can robots cook?

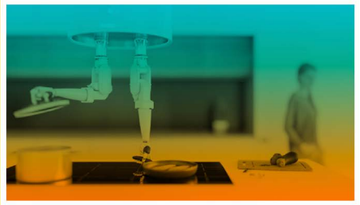

Photo credits: Canva

Alexane Jouglar, (Unamur) translate benchmark results into a language understandable by AI noobs.

Alexane Jouglar on Culinary challenges for advancing artificial intelligence

Meet the humans behind MUHAI!

Study without ChatGPT… to work more wisely with AI

Photo credits: © Pexels, modified by Folco Soffietti

Paul Van Eecke (VUB), Katrien Beuls (UNamur) and Tim Brys (VUB) explain why students should avoid the use of generative AI during high-school education to be able to use it

Outputs > Engaging policy & industry

Photo credits: © Kontrast Medien und Marketing

i2b Bremen was an important moment for the MUHAI project to present key results to the public and fellow researchers. The group from University of Bremen ensured a succesful outreach in this first final event of the MUHAI project.

Find out more

New papers have been issued, check them out!

Go to Muhai Papers repository

Project contributions > NAI webinars

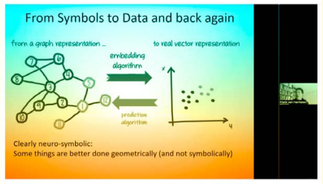

May, 2024
Frank van Harmelen, Professor in Knowledge Representation & Reasoning at Vrije Universiteit Amsterdam, gave a presentation 'The K in 'neuro-symbolic' stands for 'knowledge''. The video is available for free on youtube and is the first webinar in the NAI Webinar Series moderated by Pascal Hitzler, Editor-in-Chief of the Neurosymbolic Artificial Intelligence journal.

Find out more

Project updates > REB

June, 2024
One of the project outputs, the Recipe Execution Benchmark, has been updated with an explanatory video, you can check it out on the dedicated page in our website.

Find out more

Future events > EKAW 2024

26th-28th November, 2024
The conference serves as a premier gathering for scholars, researchers, and practitioners in the fields of knowledge engineering and knowledge management. The 24th International Conference on Knowledge Engineering and Knowledge Management (EKAW-24) will take place in Amsterdam, the Netherlands. The conference will be held at the Amsterdam Science Park campus from November 26th to November 28th 2024.

Find out more

Past events > (AKR³)

26th-30th May, 2024
The workshop at ESWC 2024 is dedicated to Knowledge Representation and Reasoning (KRR) in the area of cognitive robotics, with the focus on acquiring knowledge from the Web and making it actionable for robotic applications. Its aim is to bring together the European communities specializing in KRR and robotics to increase collaboration and accelerate advancements in the field.

Find out more

Past events > AISCICOMM24

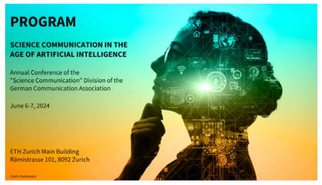

Picture: event banner

From the 6th to the 7th of June, in Zurich, researchers, communication practitioners dealing with AI gathered for an intense conference. VIU presented a study focused on the perception of AI researchers within the MUHAI project.

Find out more

Past events > ChatGPT: What now?

Picture: event at ISI

On the 9th of May 2024 Luc Steels (VIU) offered a talk at ISI Foundation (Turin, Italy) to discuss how generative AI may be improved by addressing hallucinations.

Find out more

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951846.

Unsubscribe


Click here to view it online

MUHAI Newsletter #12 | November 2024

MUHAI

Meaning and Understanding in Human-centric AI

November #wordofthemonth

AGENT – In AI, is not 007, but is like a smart helper. It can understand its surroundings (observe), make decisions (decide), and do things (act). For example, if you tell it to clean up a messy room in a game (observe), the agent would figure out how to do it (decide), and execute (act).

11/2024 MUHAI

Check out the previous words of the month!

Outputs > Volume 2

Narrative-based Understanding of Everyday Activities: An AI Cookbook

editors Luc Steels Robert Porzel

MUHAI

September, 2024
The second volume of MUHAI is accessible at this link! The volume focuses on theoretical research and concrete case studies about understanding everyday activities in the real world. The case studies come from the domain of cooking: the preparation and execution of recipes and the invention of recipes by constructing variations of existing ones.

Project Outputs > Milestone 9

October, 2024
As part of this milestone, two key knowledge graphs (KGs) were developed: the OKG includes social memories from contemporary data from twitter, and the MIRA-KG includes historical scientific data described in research articles.

Find out more

Website updates > #trAIning

#trAIning

MUHAI

October, 2024
We enriched our website with a new section to make available educational materials, especially video tutorials and webinars on aspects of AI covered by MUHAI or results developed by the project. Enjoy our brand new feature #trAIning!

Find out more


Past events > VII Frontiers Dialogue

Technological disruption: social challenges of the Artificial Intelligence revolution

VII Frontiers Scientific Dialogues

13th November, 2024
Scientific coordinator Luc Steel will participate in the VII Frontiers Scientific Dialogue titled "Technological disruption: social challenges of the Artificial Intelligence revolution" It took place at the Reial Acadèmia de Ciències i Arts, Barcelona.

Find out more

Past events > Chiang Mai Lecture

Making of MUHAI and EASE

26th-30th May, 2024
Project partner Prof Rainer Malaka from University of Bremen gave an international lecture entitled "AI - Hype, Hysteria, and Reality – Towards Human-Centric and Empowering Digital Media." Muhai was therefore presented at the Chiang Mai University, College of Arts, Media and Technology.

Find out more

Past events > MUHAI final meeting

Picture: @random stranger passing by San Zaccaria's church

From the 12th to the 14th of September Partners gathered in Venice to exchange updates on joint efforts and pave the way to the last months of the project.

Find out more

Past events > REACT

Picture: event at Digital Media Lab Bremen

In November, The Digital Media Lab Bremen recently hosted the REACT workshop, focusing on Responsible and Ethical AI in Conversational User Interface (CUI) Technologies. Organized by Nima Zargham, Thomas Mildner, and Professor Rainer Malaka, the event brought together a distinguished group of researchers from Eindhoven University of Technology, University College Dublin, and Stockholm University.

Find out more

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951846.

Chapter 4: Enlarging the community; Press Coverage

MUHAI partners were also active in engaging with traditional media, ensuring that the project's insights reached a wider audience beyond the academic sphere. Over the course of the project, twelve articles were published in different countries and languages, either authored by MUHAI partners themselves or featuring interviews with consortium members.

These articles appeared in prominent European media outlets, including:

- Il Manifesto
- Omalius
- Freie Hansestadt Bremen
- De Tijd
- De Standaard
- Knack
- Buten un binnen
- KlarText
- Nieuwsblad
- Mind
- Rai TGR Leonardo

By ensuring media outreach, MUHAI was able to bring discussions about human-centric AI and its implications into broader public discourse, bridging the gap between cutting-edge research and societal impact.

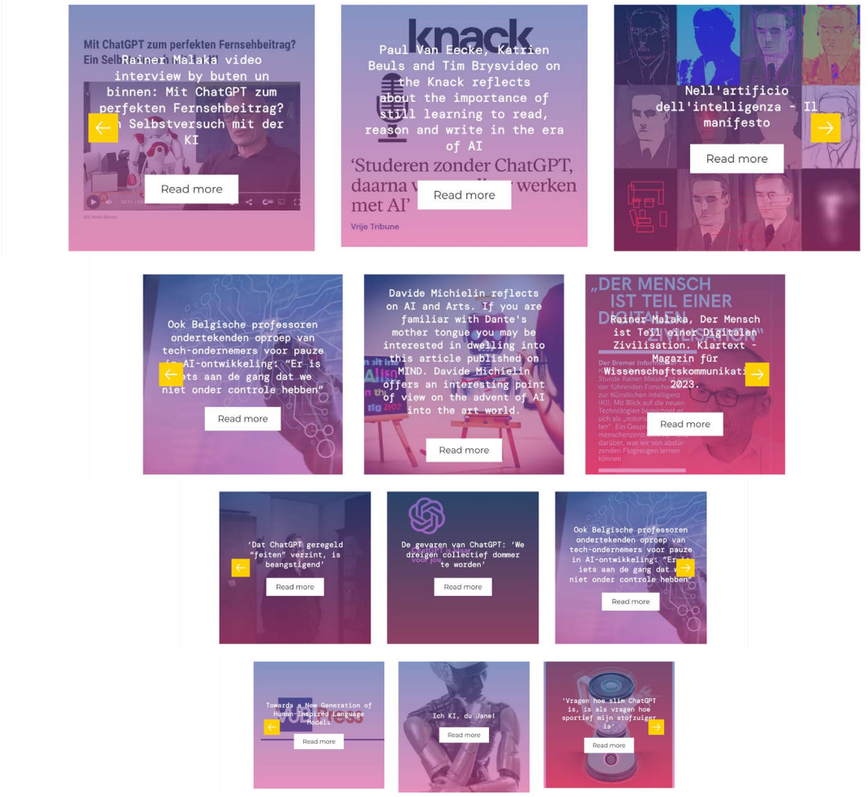
MUHAI in the Press

Mit ChatGPT zum perfekten Fernsehbeitrag? Ein Selbstversuch mit der KI

Rainer Malaka video interview by buten un binnen: Mit ChatGPT zum perfekten Fernsehbeitrag? Ein Selbstversuch mit der KI

Read more

knack

Paul Van Eecke, Katrien Beuls and Tim Brysvideo on the Knack reflects about the importance of still learning to read, reason and write in the era of AI

'Studeren zonder ChatGPT, daarna werken met AI'

Vrije Tribune

Read more

Nell'artificio dell'intelligenza - Il manifesto

Read more

Ook Belgische professoren ondertekenden oproep van tech-ondernemers voor pauze AI-ontwikkeling: "Er is iets aan de gang dat we niet onder controle hebben"

Read more

Davide Michielin reflects on AI and Arts. If you are familiar with Dante's mother tongue you may be interested in dwelling into this article published on MIND. Davide Michielin offers an interesting point of view on the advent of AI into the art world.

Read more

„DER MENSCH IST TEIL EINER

Rainer Malaka, Der Mensch ist Teil einer Digitalen Zivilisation. Klartext - Magazin für Wissenschaftskommunikation. 2023.

Read more

'Dat ChatGPT gereageld "feiten" verzint, is beangstigend'

Read more

De gevaren van ChatGPT: 'We dreigen collectief dommer te worden'

Read more

Ook Belgische professoren ondertekenden oproep van tech-ondernemers voor pauze in AI-ontwikkeling: "Er is iets aan de gang dat we niet onder controle hebben"

Read more

Towards a New Generation of Human-Centric Computer Models

Read more

Ich KI, du Jane!

Read more

'Vragen hoe alles ChatGPT is, is als vragen hoe sportief mijn stofzuiger is'

Read more

Chapter 5: Making Science Understandable and Relevant to Everyday Life: MUHAI Blogs

The MUHAI blog series was conceived as a way to translate complex research insights into accessible narratives for a wider audience. This initiative aimed to extend the impact of MUHAI beyond the scientific community, making AI-related discussions relevant to everyday life.
To ensure a balanced and structured publishing approach, the blog was organised into thematic categories reflecting MUHAI's research areas:

- **Artificial Intelligence**
- **Human-Centric AI**
- **Understanding Society**
- **Understanding Everyday Activities**
- **Art and Science**
- **Evolution of the MUHAI Blog**

The blog content evolved through three distinct phases:
Laying the groundwork – The initial phase introduced fundamental AI research themes and the broader state of the art in AI development.
Translating research into real-world applications – Subsequent entries focused on directly conveying project results, showcasing tangible insights from MUHAI's work.
Engaging with the public debate – With the rise of Generative AI and an increasing societal interest in AI's role, MUHAI blogs began contributing to broader discussions on AI's ethical and educational implications.
A notable example is the article "Study Without ChatGPT… to Work More Wisely with AI", written by Paul Van Eecke (VUB), Katrien Beuls (UNamur), and Tim Brys (VUB), and adapted by Folco Soffietti (VIU). This piece argues that students should avoid relying on Generative AI in high school, enabling them to develop a more effective and informed approach to using AI later in their academic and professional careers.

Featured Blog Entries

The following pages present a selected entry from each category, showcasing the diversity of topics covered:

Artificial Intelligence: AI at a Crossroads, by Luc Steels (VIU)
Human-Centric AI: Framing reality, by Remi van Trijp and Martina Galletti (CSL)
Understanding Society: Understanding Society, by Lise Stork (VUA)
Understanding Everyday Activities: From Kitchen to AI: A Task-Based Metric for Measuring Trust, by Robert Porzel (UHB)
Art and Science: AI, the Winning Artist?, by Folco Soffietti (VIU)

By structuring blog content in this way, MUHAI successfully fostered public engagement and ensured that its research findings were not confined to academic circles, but rather became part of the wider dialogue on AI's impact on society.

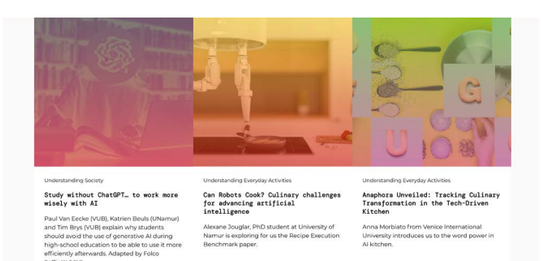

AI at a Crossroads

Luc Steels

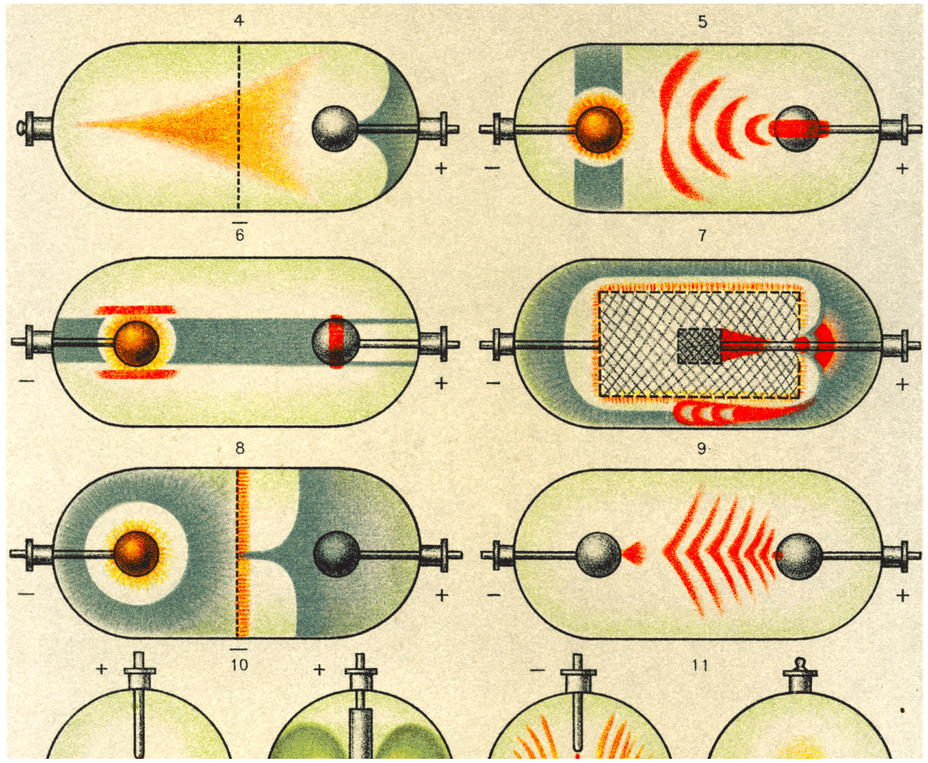

Many people believe AI (Artificial Intelligence research) started quite recently, like five years ago. But in fact the field has already had 70 years of fascinating history. It all began in the nineteen-fifties when the potential power of information technology was becoming clear, at least to a small group of far-sighed thinkers including Alan Turing and Norbert Wiener. They started to dream in earnest of building machines that exhibit some form of intelligence. Initially, there was a strong interest in constructing small cybernetic robots now called animats. For example, Grey Walter built two robotic turtles called Elmer and Elsie in 1950 and showed how they could roam around a living room, find a charging station and recharge themselves. In the same year, Claude Shannon demonstrated an electronic mouse Theseus that could learn to find its way in a maze. These animat builders were strongly influenced by the behaviorist school that dominated psychology at the time. Among the main ideas coming out of their work were models of neural networks to implement associative and reinforcement learning, which they demonstrated in behavioral conditioning scenarios similar to those used by Skinner and others to train animals.

The neural network models coming out of this work were initially applied primarily to problems in pattern recognition, image classification and systems control. In subsequent decades, neural networks and their applications multiplied, although the general idea always remained the same: Neural networks are numerical decision-makers. A network consists of several layers of nodes (loosely inspired by biological neurons) that make a weighted decision to produce a numerical output given a set of numerical inputs. For example, a node might produce a control signal to increase the speed of the left motor if a light sensor mounted on the right side of the robot captures light, so that the robot turns towards the light. More generally, neural networks

implement dynamical systems that map vectors of numbers (for example a sequence of values produced by a digital camera) to other vectors (for example a stream of signals controlling the operation of robot actuators), possibly with extra layers of decision-making in between. The intermediary layers may extract additional information from the sensory inputs or coordinate different aspects of the output.

The most remarkable property of these neural networks is that they learn autonomously based on a procedure, called a learning algorithm, that gradually changes the weights of the various decision nodes in order to minimize a decision error. Neural networks are therefore no longer programmed as is the case for ordinary computer programs. The 'trainer' only needs to provide a very large set of input-output pairs or reinforcement signals and - if all goes well - the weights then get progressively pushed in the right direction given an adequate learning algorithm.

Around 1955 a group of adventurous researchers including Herbert Simon (who later got a Nobel prize in economics), Alen Newell, Marvin Minsky, and John McCarthy opened a second thread in AI research. They focused on human mental tasks, rather than animal behavior, and started to use the term 'artificial intelligence' for their work. At first they were particularly interested in mathematical theorem proving, problem solving, board games and puzzles. By the end of the nineteen-fifties they already showed impressive demonstrations of computer programs capable of excellent performance in these domains. For example, Newell, Simon and Shaw already demonstrated around 1958 a system that could prove most of the theorems contained in the Principia Mathematica of Bertrand Russell and Norbert Whitehead.

The basic idea behind these achievements is that human intelligence is based on the creation and manipulation of symbolic structures. Symbolic structures are graphs where the nodes and links between nodes are labelled. For example, the problem of finding a path in a city is handled by representing streets, building and other landmarks as symbol nodes and the locations and spatial relations of these entities as labelled links between these nodes. Finding a path then consists in traversing this network to search connections between an initial starting point and a goal destination. Playing chess is done by representing the pieces of chess and their positions as symbols, and defining symbolically the possible moves that each piece can make on the chess board. To decide on the next move, the player generates a search space which considers the different possible moves from the current board position and evaluates whether they will give an advantage or create dangerous conditions that might lead to check mate. Because the search space of possible chess moves is very large, human players bring heuristics to bear. Heuristics are strategies to minimize search by applying more knowledge, for example, knowledge about typical openings or end-game solutions. In early AI, research into heuristics and how they could be learned was one of the main topics.

The first difference with the earlier neural network models is that this kind of AI uses symbolic representations and operations rather than vectors of numbers and numerical operations over them. It is therefore also called symbolic AI as opposed to the numerical AI of neural networks. A second difference is that this kind of AI took from the beginning the side of cognitive psychology, in opposition to the behaviourist psychology that inspired the neural network pioneers. Behaviourists argued against complex mental processing claiming that competence, even for language or problem solving, was based on fairly superficial stimulus-response associations learned through associative or reinforcement learning. Cognitive psychologists were instead no longer averse to complex internal models (like a graph of streets representing the geography of a city), rich knowledge representations (for example, semantic networks representing the common sense implications of basic concepts), or sophisticated syntactic and semantic processing (as needed in parsing and producing language). They argued instead that the stimulus-response associations implemented by neural networks were too superficial to implement reasoning, language processing or other tasks we consider

to require intelligence. In AI we similarly have a dichotomy between behaviourist AI which rests its hope on associations implemented through neural networks and cognitivist AI which works with complex symbolic representations.

A third distinction between the two schools of thought concern learning. Whereas neural network enthusiasts emphasize statistical induction, i.e. progressively generalizing from many experiences, the symbolic models primarily emphasize `learning by being told' and constructivist learning. Learning by being told means that the learner is able to comprehend instructions or advice and incorporate and use that in subsequent problem solving. Constructivist learning means that the learner uses his available knowledge to construct new distinctions or to formulate sensible hypotheses and then test them out against reality. A single exposure is then often enough to acquire a significant piece of new knowledge, contrasting with the massive amount of data that is needed to implement the statistical induction which neural networks rely on.

By the early nineteen-sixties several laboratories exploring symbolic AI had sprung up and already very significant technical advances had been made, particularly in how to handle symbolic computation. Soon many more areas of human intelligence were explored: medical diagnosis, scientific discovery, intelligent scheduling, legal argumentation, technical design, language processing, common sense reasoning, even artistic creativity. In the decades that followed, all this research lead to industrially usable expert systems assisting human experts in problem solving. They also lead to the construction of very large knowledge bases such as the knowledge graphs that underlie today's search engines, and to natural language processing technologies that could power computer-assisted translation or text editing tools.

Fast forward to more recent times.

By the beginning of the 21st century both the numerical/behaviorist AI tradition and the symbolic/cognitivist AI tradition had reached maturity. AI was no longer in the spotlight and became an accepted branch of software engineering and computer science. The field had developed a well-established set of tools and practices for building intelligent systems and they were used in a wide range of industrial and commercial applications. Meanwhile fundamental AI research continued, exploring both neural network models and symbolic methods. More fundamental research was necessary, partly because the difference between human intelligence and machine intelligence was still very significant - and it still is today.

But around 2010, a remarkable surge in the availability of data due to the deepening penetration of information technology in human activities and a considerable jump in the power of computers caused a rather sudden growth in enthusiasm for AI, specifically for the numerical AI techniques pioneered by neural network researchers, such as deep learning and convolutional networks. Earlier on these techniques were not applicable on realistic problems due to a lack of data and computing power. But now they were. The renewed enthusiasm caught on and spread rapidly throughout the world. Management consulting companies promoted (numerical) AI as the next frontier for industry and as an essential skill if companies wanted to remain competitive in today's world. Governments drew up strategic plans for AI and new start-ups and laboratories sprung up like mushrooms. The enthusiasm was not only due to the use of neural network methods. Many existing techniques of numerical and statistical analysis (such as regression, clustering, principled component analysis, optimization techniques, etc.) were now also promoted as being part of AI, thus rapidly increasing the scope of the field to encompass a far larger range of techniques and applications, beyond neural networks and symbolic methods.

But the growing reach of numerical methods and the fact that they were now labelled as AI came with a catch. The symbolic/cognitivist AI approach advocates starting from

human expertise. It tries to model human reasoning, human knowledge and human forms of communication so that the decision making by a system can be followed by a human, an explanation in human terms can be provided easily, and the system can accept advice from a human in a symbolic form (i.e. in human language). This kind of AI is therefore human-centered. It attempts to empower humans rather than replacing them. In contrast, the numerical/behaviorist AI tradition, including the recent addition of statistical numerical methods, focuses on building systems by finding the right weight parameters that (ideally) give adequate performance, but the basis of their decision-making is hidden in millions of numerical parameters that are entirely incomprehensible to a human observer, even to the designer of the network or the trainer. Such systems are forever black-boxes.

A black-box approach is alright for domains where a human-centered approach is not required, for example, for a controller of a complex technical device. But it is another matter if these numerical methods are used for domains that touch on human concerns, for example, to decide whether a prisoner gets parole, a citizen gets social housing, a consumer gets more credit, or a candidate gets an interview for a job. In those cases, the black-box approach of numerical AI becomes problematic and those who are affected by these decisions rightly feel helpless and treated unfairly. Of course numerical methods have been used for a long time (such as in operations research) but the systems built on this basis were not called intelligent. Nobody was expecting an explanation and nobody was claiming that they were as good or better than human experts. However, if you call such systems ‘intelligent’, the expectations of users increase drastically and they expect similar functions as we find in human intelligence, in particular the capacity to explain how a decision was made or to accept counterarguments, transparency, and consistency.

An additional problem of statistical methods is that they do not give the robustness and reliability that we normally expect from engineered systems. If decisions are based on statistical grounds, there are always going to be outlier cases which do not fall in the most common range. There is always going to be a bias in the data that is used for training. A decision can only be based on the features that were available for training, which might not include crucial properties of the context that a human expert would effortlessly take into account or aspects of reality which cannot be measured easily but are nevertheless important. For example, a legal advisory system built using deep learning will perform induction over a large number of cases to build statistical models how cases have been handled in the past. A new case is handled by comparing it to these models, but the system has no explicit notion of the underlying law or common custom and can therefore not justify its decisions in terms that would stand up in court. In contrast, a symbolic legal expert system will be based on a codification and implementation of the law and it will handle new cases through logical inference based on the implementation of these legal rules. This is not without its problems either, because, even in the case of codified law, there is always an interpretation step that relies on human empathy and common sense knowledge, which is very hard to capture in explicit rules.

So AI finds itself in an impasse. Numerical AI has caused great enthusiasm lately but, because it is not human-centered, it has raised a wide range of ethical and legal considerations and has generated justified worries by those caring about the rights of citizens. Particularly in Europe, this has lead to calls for developing trustworthy AI, although it is far from clear how this can be done for AI systems built by using statistical numerical methods on big data. On the other hand, we do not want to forego the obvious power that these statistical numerical methods provide either. They have proven their worth in many areas particularly in pattern recognition or systems control. So how to resolve this paradox?

My feeling is that we should do two things.

As a starter, we need to develop hybrid AI which uses both numerical approaches and symbolic approaches. Indeed, this is already happening in a number of innovative projects. For example, numerical AI is useful for learning heuristic decision rules in tandem with a symbolic system that creates search spaces using an accurate model of the domain. Numerical methods are useful for quickly retrieving information from very large knowledge bases but the knowledge-bases themselves are symbolic and the application of information to a concrete case is done with symbolic inference. Numerical techniques are effective in pattern recognition, for example for image processing and interpretation, but these techniques only give reliable results when complemented by common sense knowledge and inference to interpret the hierarchical structures and activities of a scene.

Second, fundamental AI research has to go back to the drawing board. So far both the numerical and the symbolic approach have always tried to circumvent meaning and understanding, even though meaning is central to humans as persons. A judicial decision on parole is not just a matter of statistics or the cold application of logical rules. A human judge will try to understand the social context of the offender, the motivation for the crime, the psychology and attitudes of the offender, and so on. When we send in a cv for a job, we expect that the recruiter will go beyond superficial features of a cv and build up a total picture, which includes social skills, history of achievements (even if they have nothing to do with the job itself), respect for human values, motivations, fluency in other languages, fit with other members of the team, etc. AI is not at all capable today of constructing the kind of narratives that humans make all the time in order to interpret the world and the behaviour of others. As long as that is the case, we should not throw AI into society for applications that touch on human life.

Framing reality

Remi van Trijp and Martina Galletti, CSL.

One of the reasons why it is so difficult to develop human-centric AI systems is that such systems need to “understand” the world and human activities in a way that is compatible with how humans make sense of the world. The crux of the matter is that each person has their own unique way of doing so: reality is so mind-bogglingly complex that we constantly need to make choices about which information is relevant, and which elements of a situation should be highlighted or obscured. This process – in which a person puts a situation in a particular perspective to express their beliefs, desires, and intentions – is called “framing”.

A “frame” is a structured piece of knowledge that we build up and maintain through experience. At its most basic level, a frame can be considered as a template of a scene with several open roles (called “Frame Elements”) that need to be filled in.

Figure 1 illustrates a scene that we may perceive, and everyone who has ever prepared a meal in a kitchen can immediately “frame” the scene. There are multiple frames possible, such as “Cooking”, but here we decided to frame the woman’s activity as a “Baking” event. The Baking frame includes several Frame Elements such as the person who does the baking, the food that is being prepared, utensils for doing so, a time and place (usually the kitchen), and so on. We can then communicate about what is happening depending on which aspects of the frame we wish to emphasize:
The woman is baking a cake (with the frame elements: baker + baked food
She is stirring in the pot (with the frame elements: baker + utensils)
She’s in a kitchen (with the frame elements: baker + location

Frames can also be more complex. For instance, they can propose a particular viewpoint on an event, allowing us to see the same event from different viewpoints. A classic example is the “buyer” versus “seller” frames, in which the same transaction can be viewed from the perspective of the buyer (e.g. “She bought cookies from my niece”) or the seller (e.g. “My niece sold her some cookies”). Frames can also “evoke” a wide range of associations. For instance, sentences such as “COVID-19 is an invisible enemy” evokes frames about fights and wartime, and invites the addressee to make sense of the COVID-19 pandemic in such terms. The way a society perceives a particular issue may have important effects on policy making: if COVID-19 is seen as an enemy of the people, citizens will perhaps demand more far-reaching action from their governments. If however COVID-19 is seen as “a hoax”, people might resist any new policy decision. In other words, if we want artificial systems to understand how humans perceive and make sense of complex issues, we need to identify how they “frame” particular events. Since language is one way to peek into the human mind, we can use people’s linguistic behaviors as evidence for how they frame reality. We are therefore working on a “Frame Extractor” that aims at identifying which frames people express in natural language texts written in Dutch, English, French, German, Italian and Spanish. Since these European languages are sufficiently similar to each other, we are working on a single repository of frames that can be shared by each language model.
One example is the Causation-frame, which (as you can guess from its name) frames an event in terms of a Cause-and-Effect relation. For example, in the sentence “Respondents believe that the coronavirus will cause an increase in income inequality in their country”, the Causation frame imposes a causal relation between “the coronavirus” (cause) and “an increase in income inequality in their country” (effect). The same sentence also expresses other frames: the Belief-frame (with “respondents” as the Believer, and the whole subclause as the Belief), and Future-frame (with the

auxiliary will marking that the Causation-frame is a future possibility).

So how does our Frame Extractor work? A typical workflow starts by preprocessing and preparing a document using neurostatistical language processing tools (such as SpaCy), which includes tasks such as part-of-speech tagging (e.g. recognizing whether a word is a noun or a verb), tokenization (dividing a document into sentences, and sentences into words), and dependency parsing (identifying syntactic relations between words). Using another software library we developed in the MUHAI project, the result of preprocessing is then automatically translated into a symbolic representation that is readable both by human experts as by our computational platform called Fluid Construction Grammar (FCG), which is an open source special-purpose programming language for implementing grammars based on the notion of a “construction”. A construction can be thought of as a mapping between a form or a (syntactic) pattern on the one hand, and a meaning on the other. These constructions are used for identifying which parts of a sentence can be associated to the frames in our shared repository, and which frame elements are expressed in the sentence.

As an illustration, let us translate the aforementioned coronavirus example in Italian, and see how the Italian Frame Extractor is able to identify the Causation frame: “Gli interpellati ritengono che il coronavirus causerà un aumento della disuguaglianza di reddito nel loro paese.”

Frames are typically identified by lexical or idiomatic constructions, which map a word or phrase onto a frame definition. In this example, we have a very clear lexical unit that triggers the Causation-frame: “causerà” (“will cause”). This construction will then introduce the Causation-frame, but now we still need to identify which Frame Elements are expressed. This task is performed by grammatical constructions: since the verb occurs in the Active Voice, we can infer from its lexical definition that its subject (“il coronavirus”) is the Cause, and that its Direct Object (“un aumento della disuguaglianza di reddito nel loro paese”) is the Effect. As can be seen in Figure 2, the Italian Frame Extractor indeed successfully identifies the Causation-frame and these two Frame Elements, and highlights the corresponding phrases in the text.

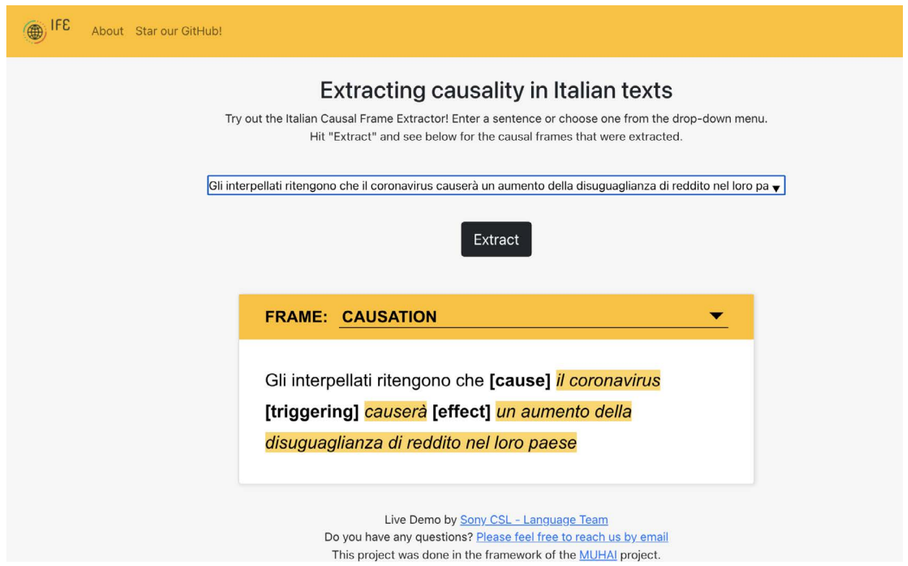

Depending on the kind of human-centric AI system that we want to develop, there are two ways of extracting frames that we call “phrase-based” and “meaning-based”. A phrase-based frame extractor can be seen as some kind of text annotator: given a

document, it needs to identify which parts of the text may evoke a frame (“Frame-Evoking Elements” or FEEs), and then annotate which phrases are assigned to which frame elements, as we showed in Figure 2. While such a Frame Extractor is relatively shallow in the sense that it does not try to comprehend a text, it is already very useful for important applications, such as extractive search (e.g. for journalists or policy makers who need to browse through large amounts of documents). Meaning-based frame extraction, on the other hand, is not about annotating a text but translating it into semantic representations that are useful for other tasks that require comprehension. Again, many applications can be envisaged for this type of frame extractor: it can be used for automatically populating ontologies and event-based knowledge graphs based on textual data, for improving difficult tasks that require more semantic information such as Entity Linking and Reference Tracking, and so on.

Since the MUHAI project is all about human-centric AI, we will make our Frame Extractors publicly available to the research community as an open source software library, with a first release for English and Italian in March 2022; followed by a yearly update and release for French (2023) and for Spanish and German (2024).

Credits
Intro photo by Visual Stories || Micheile on Unsplash
Photo of Figure 1 by Jason Briscoe on Unsplash

Understanding Society

Lise Stork

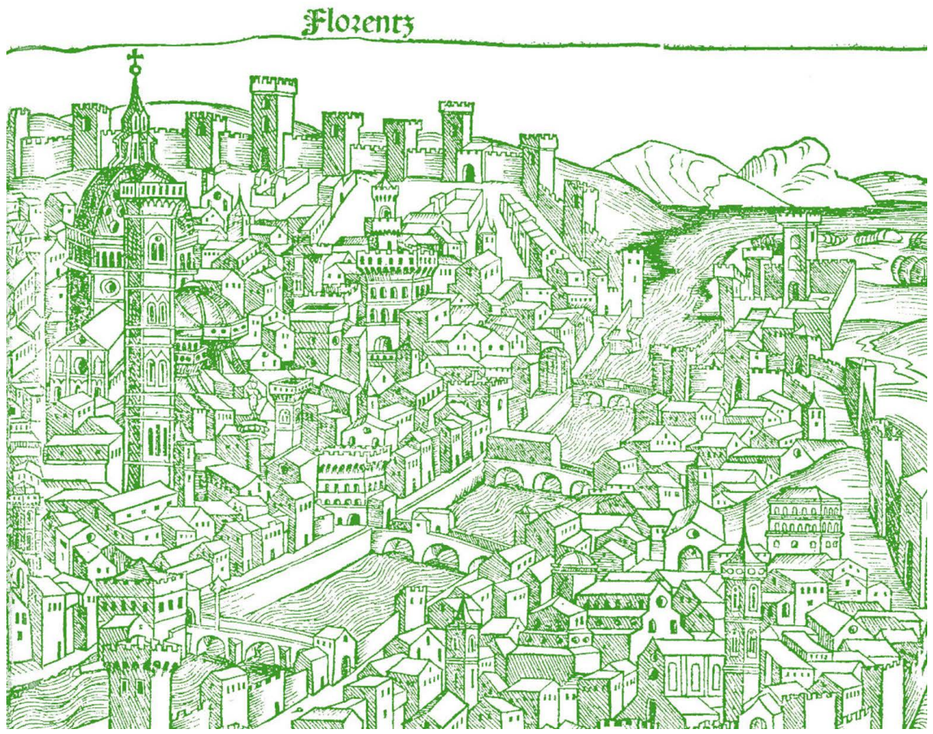

Why are the neighbourhoods in some cities sharply divided along income boundaries, while in other cities not? Was this always the case in different periods of history? And in different cultures? Has social mobility increased or decreased over time? Why does life expectancy correlate with income?

Disparities in income and opportunity for personal development are continuous sources of frustration and social divide. The deeply unequal global landscape of modern society increases the importance of studies into the origins and persistence of inequality. To this end, the MUHAI project aims to develop a technological infrastructure to aid social scientists with the generation and explanation of research hypotheses. In building such a “social observatory” the emphasis lies on cooperation between human and system, where capabilities of both complement one another.

Social scientists commonly search for indicators that contribute to or cause the origins or persistence of inequality between social groups. They do this by taking a close look at data that describe the results of societal mechanisms, such as the division of labour and income. At the Dutch International Institute of Social History (IISH) for instance, social scientists investigate the global development of labour and labour relations. For this purpose, they collect, process and link historical archives such as handwritten census, accounts of the history of municipalities, registers of births, marriages and deaths, tax surveys, and historical maps.

Source: https://stories.datalegend.net/catasto/

Research into historical datasets is challenging on multiple levels. For example: the meaning of occupations and other concepts change over time, datasets often contain biases whenever data are collected in specific regions or among certain social groups (for instance only among those that earn more than the marginal income), and although it is possible to detect certain trends computationally, finding an understandable explanation for the cause of such trends seems an insurmountable task.

Knowledge graphs are interconnected networks of data that represent historical facts and knowledge about social phenomena and everyday life. In such graphs, nodes represent real-world entities such as events, locations, or individuals of a population, and edges indicate their relationships with other entities, such as the age or birthplace of a specific individual, e.g., person -> lives In -> Paris. Researchers from the IISH routinely analyse these graphs to discover patterns and find explanations for social phenomena such as socio-economic inequality. Using different techniques that will range from statistical techniques such as deep neural networks, to symbolic techniques such as automated reasoning, we will aid them with this process by discovering new knowledge, detecting clusters or trends, and most importantly, by formulating sensible causal explanations of such clusters and trends. A typical example would be the question why in 1814, the marriage numbers in France were double that of the years before and after? The explanation for this is that Napoleon issued a law, requiring all men who were unmarried by 1815 to join the army, so a lot of marriages were hastily arranged in 1814!. Our ambition is to develop human-centered AI techniques that can uncover such explanations by working in collaboration with social scientists.

So, rather than simply uncovering statistical patterns, we aim at creating a social observatory that provides social scientists with human understandable explanations of trends, such that scientists can turn these explanations into testable hypotheses, and obtain a deeper understanding of the value of certain hypotheses and potential sources of reasoning errors, such as selection bias and missing information. The aim of this observatory is therefore not to replace human capabilities, but to enhance them, with the MUHAI tools working in collaboration with them. Through the creation of understandable narratives of social inequality, researchers can improve and accelerate their research and paint a picture of general societal processes that cause long-standing societal inequality.

The Florentine Catasto of 1427
The Catasto is a tax assessment of the inhabitants of Florence and its surrounding territories between 1427 and 1429. In the Catasto, officials listed the wealth, debts, and assets of households in the Republic of Florence. Unlike many premodern tax assessments - which only taxed the rich - the Catasto aimed to include all households within the Republic. The Catasto allows to examine the relation between household size and wealth. In the city of Florence this relation was quite positive:

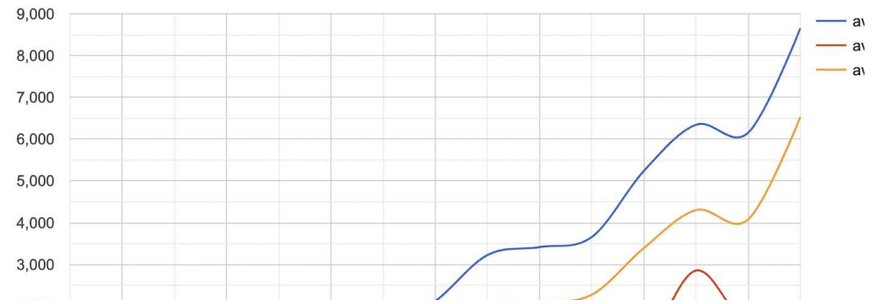

How can this correlation be explained? Were these households large because they could afford it, or were poor households also large because many family members were living under one roof?

From Kitchen to AI: A Task-based Metric for Measuring Trust

Robert Porzel

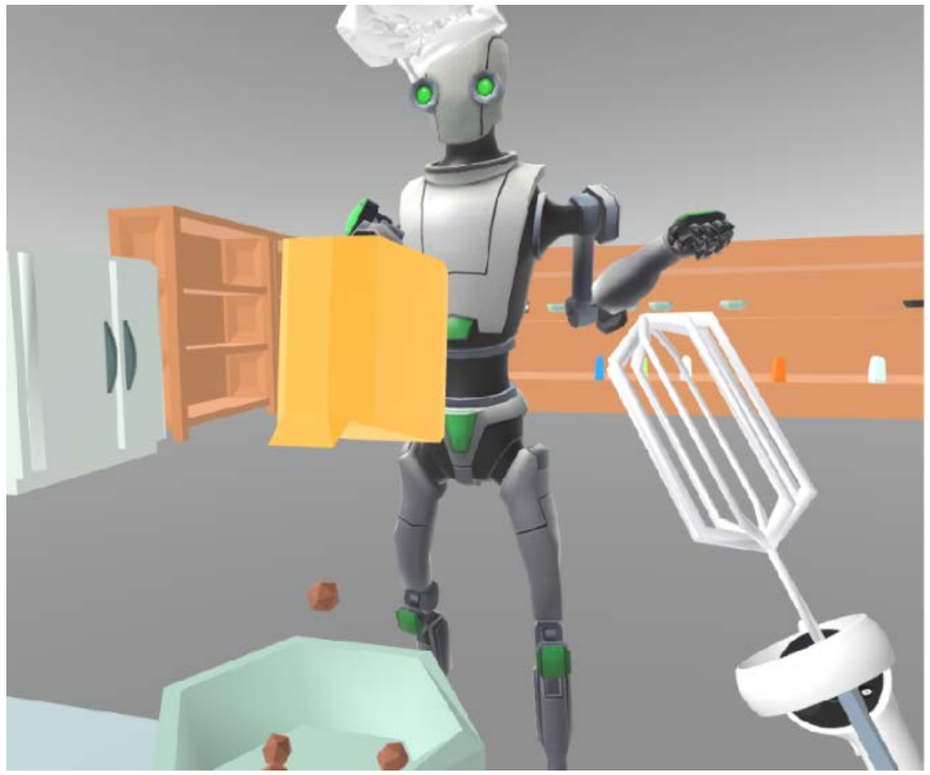

Trust is an important factor in human-centric artificial intelligence – especially for the success and effectiveness of a collaborative task in which the participants rely on each other to achieve specific sub-goals. For example, in household environments, such as a kitchen, mistakes can be made by either party that could not only lead to failure to complete the task, but even to injury through various hot or sharp appliances.
Trust in a new system or technology is critical to its success, since people tend to employ systems that they trust, and reject systems that they do not trust.

In the last few years, artificial agents, such as vacuuming robots, have become more common in household environments, and assistants for more complex tasks as cooking or cleaning are being developed. To ensure that these new systems will be accepted, it is important to explore how much people trust an autonomous system to handle these tasks, how this trust changes during use and what factors lead to an increase or decrease in trust. Toward that goal, it is important to find applicable measures for trust. For this, we propose measuring a user's trust in an artificial collaborator during cooperative cooking tasks by analysing the tasks delegated to the artificial partner during collaborative execution of a recipe.

As the delegation of tasks among humans relies on trust, we propose that the tasks given to the artificial collaborator, e.g. while preparing a meal, can supply information on the level of trust the human has in them. If the human assigns intricate, dangerous or important tasks to the artificial agent, e.g. heating or cutting an ingredient, this would indicate, that they trust this partner to complete the task successfully. Should they only delegate minor tasks to the robot – for example, wiping the counter – it indicates, that the robotic partner is only trusted to fulfil simple tasks where errors could easily mitigated. Toward the goal of measuring trust based on task delegation, three different aspects of a task that could influence a human's tendency to delegate it were chosen in our approach: difficulty, risk and possibility for error mitigation. In addition, it was deemed relevant if a human would supervise the artificial collaborator during a task or even intervene.
In addition discount factors were considered that might convince a human to assign a task to a robot even though they do not completely trust the robot, e.g. tediousness of a task or inability to complete a task themselves. These aspects were combined into a basis for a scale, that can be used to determine the level of trust the human put into the artificial partner when delegating this specific task to them.

To observe humans during cooperative cooking with an artificial partner, a VR application was developed in the Unity game engine for use with an Oculus Quest HMD. In this application the user is placed in a kitchen environment together with a virtual robot. The user can interact with various objects in the kitchen by grabbing them with either their hands or the controllers and then complete various cooking tasks by moving them in appropriate ways -- e.g., moving a whisk in circular motions through a bowl containing the different ingredients to be mixed. In addition, the user can order the robot to fulfil any of the needed cooking tasks for recipe completion -- e.g., portioning a certain amount of an ingredient into a bowl -- or some supporting such as cleaning, tidying or fetching objects for the user. For these orders a delegation-type interface is used, where the user orders the robot to fulfil a task in a declarative manner, but is not required to give details on how the task should be completed.

In the future this metric could be part of a bigger set of measures for trust specific to cooperative tasks, that includes other aspects such as the phrasing of orders given to the robot. Similarly, it could be modified for further household tasks, that could in the future be assigned to household robots. Predictions made by a graphical model based on these metrics could also be used to adjust robot behavior at runtime to calibrate trust to the appropriate level for optimal cooperation. The described test environment and scale could be used in the future to explore different robot appearances and behaviors and how they affect trust, as well as trust development over time when the human can observe the robot complete tasks successfully or make errors.

AI, the winning artist?

Folco Soffietti

The recent discussion sparked around art prizes being assigned to artworks generated by AI - whether stated or only revealed afterwards - leaves space for interesting reflections.

Boris Eldagsen, Berlin-based photo & video artist, purposely entered the Sony world photography with an AI-generated work, in order to provoke a debate. Eventually he refused the Sony World Photography Awards 2023 prize considering it a flaw in the organization of the contest, since AI images and photos – he stated – should not compete with each other, and this, despite the jury having inquired beforehand if the photography had been co-created by AI. According to the jurors, it was important to acknowledge the need to address the role of AI in art creation and the broader consequences of this new practice(Grierson, 2023).

And indeed, the fact that the most recent release of Photoshop, the world leading software for photo retouching from Adobe, has embedded an AI generative tool able to modify and fill pictures, makes it urgent the need to understand and address AI in creative processes.

The world of photography is however not the only one affected, in our fast changing society: digital art can also be considered a traditional medium as proved by Jason Allen's win at the Colorado State Fair's digital arts competition with his work "Théâtre D'opéra Spatial". Several digital illustrators were not happy with this result when it emerged that the AI platform 'Midjourney' was involved in the process, and this despite the claim that 80 hours were needed to complete the artwork.
The fair's submission guidelines do not directly mention AI-generated art, but they define digital arts as "artistic practice that uses digital technology as part of the creative or presentation process", hence AI could enter. Also, the judges said "they awarded the top prize based on the story that Théâtre D'opéra Spatial tells, as well as the spirit it invokes" and this would have been the result in spite of the media involved, as far as they are digital (Kuta, 2022).

The concern of professionals and early career artists are understandable (readers will remember the anti-AI posts that appeared consistently on Instagram accounts at the end of 2022). Contests are one of the first steps to enter the art market and gain recognition. AI implies a democratization of the possibility to create artworks in several styles, virtually allowing everyone to be an artist or, at least, a content creator. This also opens the Pandora's box of copyright infringement, an issue not-so easy to address, in which the law is generally still trying to bind the copyright to a human intellectual creative process, with the notable exception of UK as well explained by Borg et Al., 2023.

It will also be interesting to discover how the future art critics and citizens will look at this early age of AI art. Will the human author be remembered or will the AI process be remembered? Or, as it happened in the past, the AI and the authors will, to the non-expert eye, be fused in an art movement or period?
Coming back to art contests, in more practical and urgent terms what should be addressed is the assessment method of the contest itself. It is likely that AI art will be inscribed into dedicated categories, even though Visual Arts categories are usually extremely broad and could allocate AI artworks. Of course it remains difficult to evaluate the level of AI involvement: is it just for retouching, is it for improved resolution or light management, is it only part of the image, is it the starting point? Regarding the latter case, one we might call "AI based", what the assessment may still consider, at least in the current AI tools, is another specificity: the prompt.
The prompt contains the instructions and, as such, reveals the intention of the human creator, hence testifying the creativity and the consistency in the idea-result. The references employed would also be presented, helping verifying that no copyright infringement occurred. Moreover, the writing style can be another indicator of quality, for instance achieving a certain visual result with a poem, a rich lexicon, or an effective sentence.

The full implication of AI to the art work are only at their dawn, and if curiosity on how art history will evolve is high, also regarding what machines will do, there is currently the need to ensure that art contests are transparent, that traditional and AI-related entries have categories and opportunities to be assessed and awarded. Transparency of rules and clarity of indicators will help the current and future generation of artists to adapt and integrate, if they wish, AI in their artistic process, without interfering with traditional processes that will coexist along AI art.

References

Borg, W. P.-J., Podoprikhina, G., & Alexander, L. (2023, February 21). AI-Generated Art: Copyright Implications. Lexology. https://www.lexology.com/library/detail.aspx?g=fff221ab-66ce-4cbb-8f0d-687850888696
Kuta, S. (2022, Sept 06). Art Made With Artificial Intelligence Wins at State Fair. Smithsonian Magazine: https://www.smithsonianmag.com/smart-news/artificial-intelligence-art-wins-colorado-state-fair-180980703/
Grierson, J. (2023, Apr 17). Photographer admits prize-winning image was AI-generated. The Guardian: https://www.theguardian.com/technology/2023/apr/17/photographer-admits-prize-winning-image-was-ai-generated

Understanding Everyday Activities

Narrative Objects

A recurring theme of AI research has turned out to be that what should be easy often is not. Consider this question – "what can I cut a stick of butter with?" You probably already thought of an answer, and if pressed, you could invent more creative ones. A string might do, or the edge of a glass perhaps if no knife is available – though you might protest if I suggested a jar's edge. It will be annoying to get the butter rests out of the threading.

Understanding Society

A Digital Assistant for Scientific Discovery in the Social Sciences and Humanities

Scientific discovery aims to explain the mechanisms that govern our world. In the social sciences and humanities, scientists are interested in our social world; societies and the individuals within them. They research, for instance, the mechanisms that cause social divides. Why do some groups die younger than others? Why do women earn less than men? How do the occupations of your (grand)parents influence your own career?

Understanding Society

Narrativizing Knowledge Graphs

Any natural language expression of a set of facts — that can be represented as a knowledge graph — will more or less overtly assume a specific perspective on these facts. In this work we see theconversion of a given knowledge graph into natural language as the construction of a narrativeabout the assertions made by the knowledge graph.

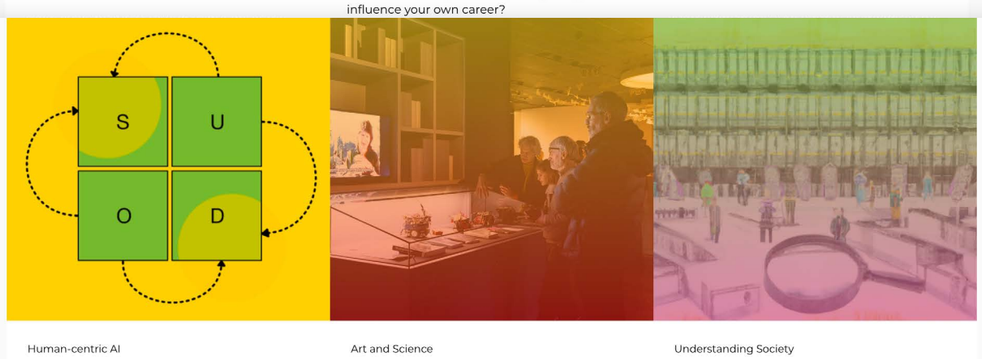

Human-centric AI

Uncommon Ground

Robert Porzel from University of Bremen, reflects on disagreement between users and AI in their respective beliefs and knowledge.

Art and Science

Science on the edge of chaos

MUHAI' scientific advisor Luc Steels guides us in the visit of a pop-up exhibition about the history of AI he curated in December, not to be missed!

Understanding Society

From digital archives to online observatories, the peaks and chasms of social-media based research Pt.3

Social media have a societal value, working as modern agoras but there is a need for understanding public concerns and the perceptions of anthropogenic phenomena at a time when user data is increasingly monetized. Part 3 of 3.

Understanding Society

From digital archives to online observatories, the peaks and chasms of social-media based research Pt.2

Social media have a societal value, working as modern agoras but there is a need for understanding public concerns and the perceptions of anthropogenic phenomena at a time when user data is increasingly monetized. Part 2 of 3.

Understanding Society

From digital archives to online observatories, the peaks and chasms of social-media based research Pt.1

Social media have a societal value, working as modern agoras but there is a need for understanding public concerns and the perceptions of anthropogenic phenomena at a time when user data is increasingly monetized. Part 1 of 3.

Art and Science

AI, the winning artist?

The fact that AI-generated entries have won art and photo contests sparked a debate in the art world. The blog reflects on what are the possible paths for art prizes in a world in which AI generators are commonly used.

Human-centric AI

Framing reality

One of the reasons why it is so difficult to develop human-centric AI systems is that such systems need to "understand" the world and human activities in a way that is compatible with how humans make sense of the world.

Understanding Society

Economists' inequality narratives (on Twitter) before and after the COVID-19 outbreak

Inequality-related narratives can be created, circulated and employed at two distinct but intertwined debate levels. The first one is made by scholarly and scientific debates, which are mostly carried out by academic researchers and field experts, specialised in the measurement, analysis and modeling of specific types of inequalities, like inequalities of access to health and care services.

Understanding Everyday Activities

Toward a formal theory of narratives

The activities of people as well as of artificial agents in reality, virtual reality or simulation can be recorded as data that discretize trajectories of body parts and the ensuing force events. While these data provide vast amounts of information they are, by themselves, meaningless.

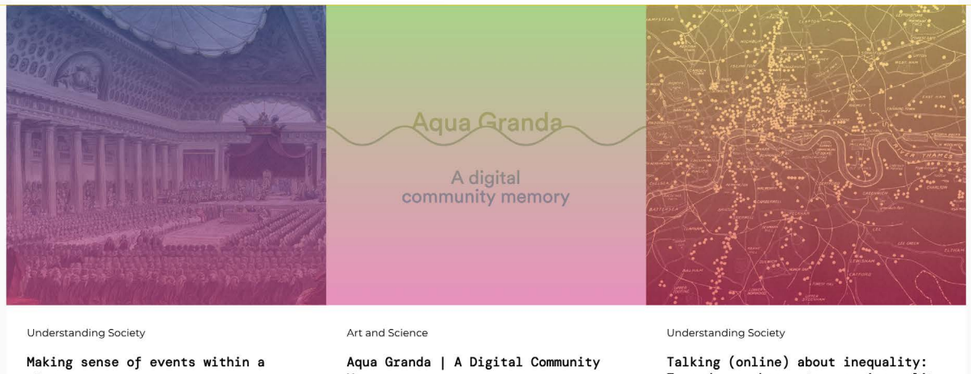

Understanding Society

Making sense of events within a story

We are constantly building a posteriori stories about how events happened, and about how you can make connections to come up with a coherent whole. To some extent, it is an extension of the Five W's - Who? What? When? Where? Why? -, the five questions that are considered the most basic in problem solving.

Art and Science

Aqua Granda | A Digital Community Memory

Within the framework of "Aqua Granda, a digital community memory" initiative, launched in November 2020 by the EU H2020 ODYCCEUS project and Science Gallery Venice, among the main outputs is a book, available in open access.

Understanding Society

Talking (online) about inequality: Towards an observatory on inequality narratives

Part of the MUHAI objectives is developing tools to help humans understand media materials, such as tweets or articles, on critical social issues, in particular socio-economic inequality.

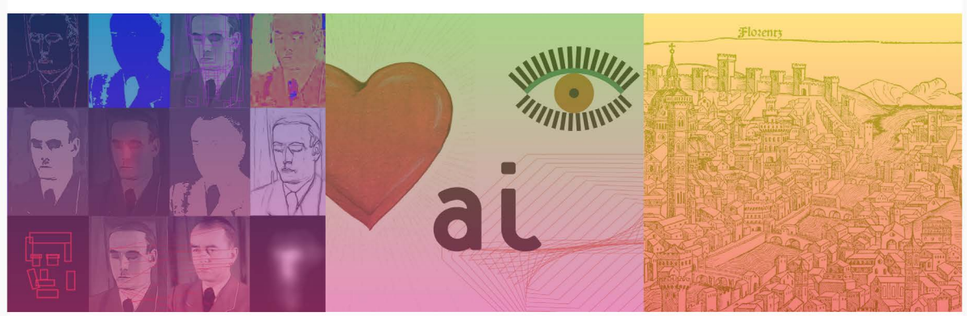

Art and Science

Luc Tuymans through the lens of AI

AI (Artificial Intelligence) researchers try to understand the structures and processes that underlie intelligence and use that insight to build practical applications. Much has already been achieved. But much remains to be discovered.

Human-centric AI

MUHAI Visual Identity

If the term "scientific" usually attracts our attention, the term "artificial" often alerts us, making us think of a world where human beings are at the margin and robots at the center.

Understanding Society

Understanding Society

Why are the neighbourhoods in some cities sharply divided along income boundaries, while in other cities not? Was this always the case in different periods of history? And in different cultures? Has social mobility increased or decreased over time? Why does life expectancy correlate with income?

Human-centric AI

Deconstructing Recipes

What is the secret ingredient of recipes? In recipes, we talk about ingredients, sometimes many of them, which get moved into different containers and transformed in a thousand different ways, thus turning into other things (such as a dough, or a puree). Moreover, these ingredients and 'resultant objects' are often unmentioned, as in 'Bake until crispy and golden.' 'Pour until saturated'. Yet, we are always capable of understanding what that specific step or instruction is talking about. How?

Human-centric AI

Foundations for Meaning and Understanding in Human-centric AI

Through the Foundations for Meaning and Understanding in Human-centric AI, the MUHAI project offers an in-depth and integrated overview of narratives and understanding in different disciplines and research fields. The volume builds upon recent insights and findings from social and cognitive sciences, humanities and other fields for which narratives have been found to play a relevant role in human understanding and decision-making processes. This, to map the state-of-the-art of narrative-centric studies and to identify the most promising research streams for tomorrow's AI.

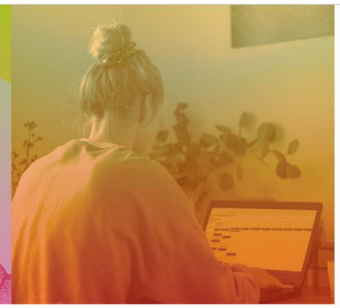

Human-centric AI

The FCG Editor: a new milestone for linguistics and human-centric AI

When people hear the word "grammar", most of them still think about a set of syntactic rules to combine words (and their concepts) in a compositional fashion. Most NLP systems therefore consider grammar to be equal to syntactic parsing, so syntax is simply one of the components of a traditional pipeline that can be useful for downstream tasks. In the MUHAI project, we take a different approach inspired by cognitive-functional linguistics, in which grammar itself is meaningful: it expresses how people conceptualise reality, and which perspective they take on the events that they perceive in their daily lives.

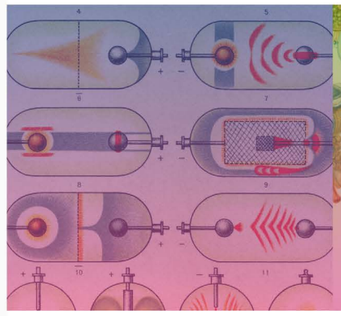

Artificial Intelligence

AI at a Crossroads

Many people believe AI (Artificial Intelligence research) started quite recently, like five years ago. But in fact the field has already had 70 years of fascinating history.

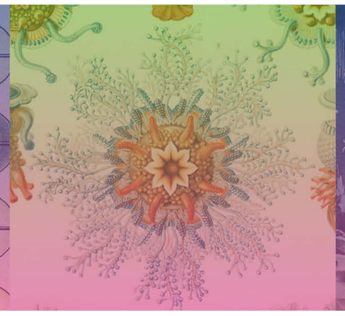

Understanding Everyday Activities

Understanding Everyday Activities

If the proof of the pudding is in the eating then the ultimate test for understanding an instruction is its proper execution. This view greatly expands the scope of natural language understanding

Human-centric AI

Visions of the Future

It started in the summer of 2030. Cecile's daughter Lucilla had visited her mother's wonderful home in a small town in the south of France and noticed that she was occasionally very absent-minded, forgetting things and behaving in a disoriented and erratic way, even for routine every day household tasks like cooking. Lucilla found a stove turned on without anything on it, several cups of half-drunk coffee around the living room, a pot with raw eggs and sugar. Her brother Juha had noticed similar unusual situations, including a chicken in the oven that had been roasted until black.

Understanding Everyday Activities

Deep Understanding of Everyday Activity Commands

Performing household activities such as cooking and cleaning have, until recently, been the exclusive provenance of human participants. However, the development of robotic agents that can perform different tasks of increasing complexity is slowly changing this state of affairs, creating new opportunities in the domain of household robotics. Most commonly, any robot activity starts with the robot receiving directions or commands for that specific activity. Today, this is mostly done via programming languages or pre-defined user interfaces, but this changes rapidly.

Human-centric AI

Linguistic Alignment for Chatbots

Over the last decade, conversational agents have slowly but surely become very common in our daily lives. Conversational agents - these are computer agents, robots or software which can understand and talk to a user in natural language - include assistants on our smartphones like Siri or Google Assistant, voice-controlled smart home devices like Google Home and Amazon Alexa, as well as online chatbots like the ones frequently used in customer service.

Understanding Everyday Activities

Curiosity-Driven Exploration of Pouring Liquids

Babies, puppies, kittens may be bundles of joy but they are also agents of pure chaos. They knock things over, stick their fingers where they're not supposed to, and get a taste or sniff of anything they can, all just for fun of course. Or is it "just for fun"? In the moment, it is, but this playfulness of infants also allows them to build the intuitive models of the physical world which are needed to cope with that world "seriously".

Understanding Everyday Activities

From Kitchen to AI: A Task-based Metric for Measuring Trust

Trust is an important factor in human-centric artificial intelligence – especially for the success and effectiveness of a collaborative task in which the participants rely on each other to achieve specific sub-goals. For example, in household environments, such as a kitchen, mistakes can be made by either party that could not only lead to failure to complete the task, but even to injury through various hot or sharp appliances. Trust in a new system or technology is critical to its success, since people tend to employ systems that they trust, and reject systems that they do not trust.

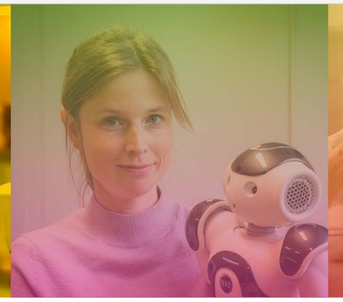

Human-centric AI

Do you speak AI?

We are pleased to announce that MUHAI researcher Katrien Beuls is now Lecturer in Artificial Intelligence at the University of Namur (Belgium) and we are happy to welcome UNamur in the MUHAI consortium! Read her interview for the University of Namur's magazine "Omalius". ©UNamur (Photo credits: ©Christophe Danaux)

Human-centric AI

Pragmatics: the secret ingredient

Imagine you're at a house party. You've just met someone new and you want to make a good impression. So, you ask them "How's it going?". To a machine, this phrase might seem like a simple question, but to a human it has a whole range of meanings. It could be a polite greeting, a genuine inquiry about how the person is doing, or even a subtle way of asking them to leave. The machine might not be able to pick up on the nuances of the conversation, but a human would be able to decipher the true meaning behind the words. That's pragmatics in action!

Chapter 6: Going to the Public: Art and AI

A significant application of human-centric AI is found in contemporary art. The intersection of art and AI was a key area of exploration within MUHAI, both as a subject of artistic exhibitions and as a powerful tool for archiving, data management, and curation. MUHAI investigated how AI could inform, inspire, and shape artistic expression while also preserving cultural heritage.
One notable example is AquaGranda, a project reflecting on the role of digital technology in memory preservation:
“Paradoxically, the very digital technologies that have contributed to making our world more fragmented and obsessed with the present may also help us create, exchange, and preserve cultural memories. This project emerged from that idea: Why not use the vast number of social media messages shared in reaction to the Aqua Granda flood of November 2019, enriching them with oral histories and historical documents from both 2019 and 1966, to create a collective digital memory? Why not make this freely available to scholars and citizens to understand the social impact of this devastating event and prepare for the future? Why not encourage the creation of artworks inspired by this collective digital memory as a memorial to trauma and loss?”
(Steels & Sartoris, 2021)
MUHAI engaged with the art world and the broader public through collaborations with contemporary artists, citizen-engagement initiatives (e.g., the AquaGranda Digital Archive), and partnerships with EU projects such as Odysseus and AI4EU, as well as with major art institutions like BOZAR and Science Gallery Venice.

Key Exhibitions

SECRETS (3 April – 2 May 2021, BOZAR, Brussels)
SECRETS emerged from the “scientist-in-residence” programme, bringing together Luc Steels (VIU, MUHAI Scientific Coordinator) and Luc Tuymans, a world-renowned Belgian painter. The exhibition explored the potential of AI as a tool for contemporary art creation, examining the differences and analogies between how AI and humans perceive paintings.
The central theme of *SECRETS*—meaning—is a fundamental concept in MUHAI’s research. By juxtaposing AI-generated interpretations with human artistic expression, the exhibition raised critical questions about the role of artificial intelligence in artistic creativity and human perception.
AquaGranda: A Digital Community Memory (Virtual Exhibition and Publication)
The devastating high tide of November 2019 in Venice inspired the creation of a digital archive, collecting both quantitative data and first-hand testimonies from those affected. This archive was transformed into a virtual exhibition, showcasing AI’s role in analysing and interpreting the event, while drawing attention to Venice’s vulnerability to climate change.

Luc Steels (VIU, MUHAI Scientific Coordinator) contributed to the creation of the archive and curated the accompanying publication, which formed the basis of a joint workshop organised by MUHAI and the H2020 project AI4EU. The exhibition was recognised with an Honorary Mention for the European Union Award for Citizen Science 2023.

Science on the Edge of Chaos (18–22 December 2023, Royal Library, Brussels)
This pop-up exhibition explored the scientific revolution of the 1980s and 1990s, when researchers across multiple disciplines recognised that complex natural phenomena—from tornado formations to ant colony behaviour—could be understood as self-organising dynamical systems. The exhibition examined how this shift influenced fields such as physics, chemistry, neuroscience, and economics, and ultimately played a role in reshaping AI research.

Curated by Luc Steels, the exhibition featured:
Archival video interviews with pioneering scientists, including Nobel laureates Christian de Duve, Manfred Eigen, Roger Penrose, and Ilya Prigogine, as well as the founders of chaos theory (Benoît Mandelbrot, Daniel Ruelle, Floris Takens, etc.).
Books and experimental objects that captured the era’s groundbreaking discoveries.
Artworks from Anne Marie Maes’ “Wunderkammer”, drawing connections between science, nature, and art.
Robots and video experiments from the 1990s, illustrating the rise of self-organising artificial intelligence and neural networks.
The exhibition also incorporated historical scientific texts from the Royal Library’s collection, including Newton’s Principia, Copernicus’ celestial observations, Laplace’s writings on planetary motion, and 17th-century botanical illustrations. This fusion of past and present scientific knowledge underscored the ongoing dialogue between mathematics, nature, and artificial intelligence.
Additionally, the project provided an opportunity to reflect on the impact of generative AI in artistic creation, a theme explored in a dedicated MUHAI blog entry.

In the picture: Science on the Edge of Cahos Exhibition - Photo by Thierry Geenen.

Chapter 7: Engaging in the Public Debate: Ethics

The rise of artificial intelligence has sparked global discussions on ethics, trust, and responsibility, and MUHAI partners have actively contributed to these conversations. From participating in dedicated conferences to publishing thought-provoking papers, the project has explored the ethical implications of AI and its role in human-machine collaboration.

One key event was the Ethics and Trust in Human-AI Collaboration: Socio-Technical Approaches workshop (21 August 2023, Macao), where MUHAI was represented by the University of Bremen. Organised by institutions including Union College, University of Brescia, Tulane University, IBM Research, and the University of South Carolina, the workshop brought together experts to examine trust and ethical concerns in AI-driven interactions. As AI becomes increasingly integrated into creative and decision-making processes—particularly with advances in generative AI—ensuring human trust remains a central challenge. Discussions focused on risks such as deskilling, displacement of human decision-makers, and value misalignment, highlighting the need for ethical AI design from the outset.

Ethical AI was also a focus at Shifting Tides: Norms, Ethics, and Values in an AI-Infused Society, a colloquium held at the University of Bologna (July 2023). In September 2023, a Venice International University (VIU) representative contributed to the debate with a talk titled Is Value-Aware AI Possible? Meanwhile, during ECAI 2023, VIU co-chaired a workshop on Value Engineering in AI, with its proceedings—co-edited by Luc Steels—set for publication in Springer's Lecture Notes in AI series. Steels' paper, Values, Norms, and AI, explores the intersection of ethics and artificial intelligence within this volume.

Beyond academic conferences, MUHAI researchers also engaged in cross-disciplinary discussions. The REACT workshop on Responsible and Ethical AI in Conversational User Interfaces (CUI Technologies), hosted by the Digital Media Lab Bremen, gathered researchers from Eindhoven University of Technology, University College Dublin, and Stockholm University. With keynotes by Benjamin Cowan, Minha Lee, and Robert Porzel, the event emphasised transdisciplinary collaboration in addressing ethical challenges in AI.

Further contributions came during VECOMP24, where Luc Steels served as Programme Chair, curating a track on values and ethics in AI. In November 2024, researchers from VUB, VIU, and Sony participated in an event on the limitations of generative AI, critically examining the risks and potential misuses of AI-generated content.
Shaping the Ethical Discourse Through Research and Writing
MUHAI's commitment to ethical AI extends beyond conferences—it has also produced research that shapes the broader public discourse. Scholars affiliated with the project have published influential papers, including:
Laura Spillner, Rachel Ringe, Robert Porzel, and Rainer Malaka (2023). "My, My, How Can I Resist You?"—Examining User Reactions to Bogus Explanations of AI. Presented at the Workshop on Ethics and Trust in Human-AI Collaboration.
Rachel Ringe and Robert Porzel (2023). Towards a Task-Based Metric for Measuring Trust in Autonomous Robots for Everyday Activities. Published in the CHI TRAIT Workshop on Trust and Reliance in AI-Assisted Tasks.

Beyond academic publications, MUHAI researchers have contributed to public debate through mainstream media. An article titled Studeren zonder ChatGPT, daarna verstandiger werken met AI (Study Without ChatGPT, Then Work More Wisely with AI), published in Knack, reflects on the responsible use of generative AI in education. The piece, authored by Paul Van Eecke (VUB AI Lab), Katrien Beuls (UNamur), and Tim Brys (VUB AI Lab), argues that students should first develop critical thinking skills before relying on AI tools, ensuring they use these technologies wisely in the future. By engaging with academics, policymakers, and the public, MUHAI has positioned itself at the forefront of the ethical AI debate, helping to shape the conversation around trust, values, and the responsible deployment of AI in society.

Conclusions: Lessons from MUHAI and Future Directions

The MUHAI project provided valuable insights into the challenges and opportunities of AI communication. As AI technologies continue to evolve, fostering trust and transparency remains paramount. The study highlights three key recommendations for future AI communication strategies:

Training AI researchers in communication: Equipping researchers with better outreach skills ensures greater public engagement and informed discourse.

Enhancing collaboration between researchers and communicators: Dedicated science communicators can help bridge knowledge gaps and improve outreach efforts.

Diversifying communication formats: A combination of traditional media, digital campaigns, and artistic collaborations can enhance AI's public understanding and trust.

While MUHAI did not dramatically alter AI discourse, it reinforced the importance of well-structured communication strategies in shaping public perceptions and policymaking. In an era where AI narratives fluctuate between utopian and dystopian extremes, projects like MUHAI highlight the need for balanced, ethical, and transparent discussions—ensuring AI serves society responsibly and inclusively.

This project has received funding from the European Pathfinder Project under Grant Agreement N° 951846