Trustworthy AI

Trustworthy AI Project Newsletter – Issue 4, August 2022

CONTENTS

Co-funded by the Erasmus+ Programme of the European Union

Our APPROACH

Integrating AI with Ethics and Trust

Trustworthy AI is a pioneering project that integrates the teaching ethics and trust into the AI curricula, following the EU High-Level Expert Group guidelines about the 7 elements of trustworthy AI.

Transversal Teaching of AI

The project brings added value by raising awareness, for the first time, of the potential, opportunities and risks of AI amongst teachers and students of all backgrounds.

Innovative Teaching Strategies

Trustworthy AI makes a significant contribution in enhancing commitment and capacity of HEIs to innovate in their teaching, not only from the content perspective, but also with regard to the methodologies.

Co-funded by the Erasmus+ Programme of the European Union

Our OBJECTIVES

Produce 3 new resources to enhance the capacity of HEIs to introduce trustworthy AI teaching in their curricula.

Rigorously test the resources with more than 48 teachers and 200 students to optimise their relevance and effectiveness.

Strategically disseminate the resources produced, reaching at least 240 teachers that will integrate the latter in their teaching.

Trustworthy AI

Co-funded by the Erasmus+ Programme of the European Union

THE CONSORTIUM

The Trustworthy AI project unites 7 partners from universities, businesses, start-ups, and networks from 5 EU Member States, whose experience and expertise provide an ideal foundation to achieve the project's objectives.

University of Alcalá – Project Coordinator
Madrid, Spain

National University of Ireland Maynooth
Maynooth, Ireland

European E-learning Institute
Copenhagen, Denmark

Umeå University
Umeå, Sweden

Stichting ALLAI Nederland
Amsterdam, The Netherlands

University Industry Innovation Network
Amsterdam, The Netherlands

Momentum Consulting
Leitrim, Ireland

Co-funded by the Erasmus+ Programme of the European Union

WELCOME TO OUR FINAL NEWSLETTER

After 2 years, the Trustworthy AI project is coming to an end. Despite dealing with a topic that does not receive as much attention as it should, the Trustworthy AI project has succeeded in raising awareness about the immense benefits as well as risks associated with the development of artificial intelligence systems.

Using the European Commission's Ethics guidelines for trustworthy AI as a guideline, the Trustworthy AI project set out with the aim of introducing a new methodology for AI teaching which emphasises, amongst other things, ethical and moral principles of artificial intelligence. From research reports, and guides to multiplier events and hackathons, the Trustworthy AI project has produced invaluable outputs that bring ethical AI development to force, making it accessible to all. All outputs created by out project are free to use.

In this edition of the Trustworthy AI newsletter, we bring to you the impact that Trustworthy AI has had. Our national and international launch events, which took place in July and August. The partners' experience in delivering the Artificial Intelligence hackathons, as well as projections on the sustainability of the resources created in the long run. The effect of Trustworthy AI can be seen in the number of event attendees, the positive feedback received from relevant stakeholders, engagement in our resources, and the testimonials from our partners on the impact of the Trustworthy AI project.

On behalf of the consortium, UIIN wish you a last pleasant read!

Photo by James A. Molnar on Unsplash

Co-funded by the Erasmus+ Programme of the European Union

News from the project

Photo by Stefan Cosma on Unsplash

Trustworthy AI

Co-funded by the Erasmus+ Programme of the European Union

MULTI-DAY HACKATHON

Trustworthy AI Hackathon at the University of Alcala

Between the 15th and the 17th of May, UAH held the first Trustworthy AI hackathon, organized using the Guide to Conducting Ethical AI Hackathons, which is the result of the third intellectual output of the Trustworthy AI project. Due to the circumstances, the event was held online, and the process and results are still available on the Devpost website that was set up for the occasion (https://trustworthy-ai.devpost.com/). The event was followed by 36 participants who enjoyed the Hackathon while being introduced to the concept of Trustworthy AI using the resources created by our project.

The event was separated between three days. On the first day, participants were introduced to the general topic of Ethics, AI and trustworthiness, which most of them found helpful as their backgrounds were varied and mostly outside the field. The participants came from various degrees and Master's courses, ranging from the technical ones (such as Data Science or Business Analytics) to less technical fields, such as Economics, Business Administration or even Criminology. They were all presented with the IO2 resources, mainly the explainer videos and the deck card, which were both used for educative purposes and for participants to get to know each other in the group formation phase. Their questions after day one were answered by the mentors (formed by a group of teachers) at the end of day one.

Day two started with the actual detailed problem. They were introduced to the Inclusive AI issues with a webinar offered by OdiseIA, with whom we partnered with; they are a Spanish NGO focused on Ethics and AI and they are formed by a large group of professionals. Students worked all day in their proposal for a particular problem and submitted their presentation at the end of the day.

The final day was the one were the jury watched the presentations and read the proposals; with their verdict, a winner was chosen – the winning proposal, about inclusive AI for e-commerce, is available at the Devpost website.

Overall, the experience was very good, and the participants opinions reflected this. As always, however, we found room for improvement; some participants, for instance, told us that three days might be too short as, in May, they are also studying for their final exams and time is limited, so a week could have been better to produce better results. This was included in the Guide as participant's feedback. We also took advantage of the occasion to test the IO2 products and, overall, the audience was very satisfied with them. The detailed results can be found on the Trustworthy AI project's website (https://www.trustworthyaiproject.eu/).

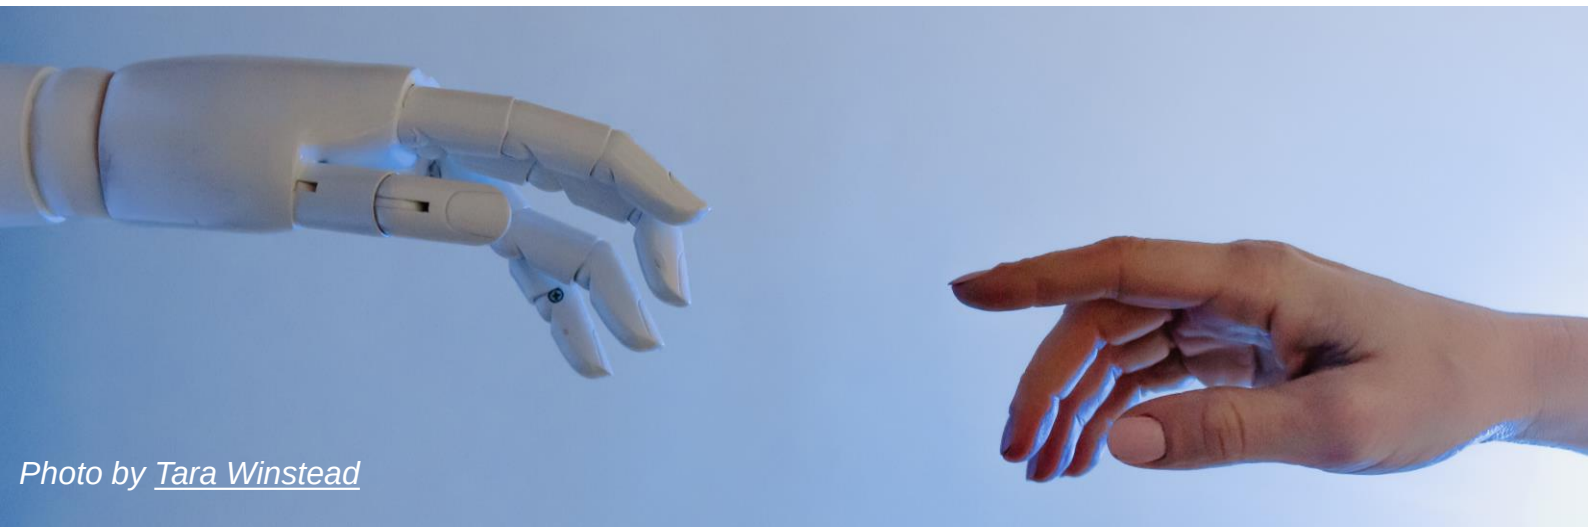

Photo by Tara Winstead

Co-funded by the Erasmus+ Programme of the European Union

Bringing Trustworthy AI to the people

Trustworthy AI Multiplier event at Maynooth University

Maynooth University, Ireland organized a multiplier event for the Trustworthy AI Project at Dublin City University on 28th July. This event aimed to present the Trustworthy AI project and its achieved intellectual objectives; followed by lunch for the physical attendees. The event started off with an introductory note from Dr Markus Helfert, the director of Innovation Value Institute (the research hub for digital innovation in Maynooth University) and Dr. Zohreh Pourzolfaghar. The three-hour event aimed to highlight three aspects of the project. First, the presentation of the project and its achievements to deepen the understanding of the relevance and importance of trustworthy AI for current and future. Second, the presentation of the very practical use-case of CHAPE (http://cohealth.ivi.ie/chape/) in the healthcare domain by Dr. Marco Alfano was used at the forefront of the Ethical AI Hackathon conducted by Maynooth University in June. This use case helped consolidate the students' learnings about the ethical dilemmas in the healthcare industry. Third, the presentation of the user testing of the intellectual outputs conducted by Maynooth University. Roshmita Kanungoe from highlighted the results of the user testing, which helped measure the effectiveness of the resources developed during the project.

Events such as these are paramount in spreading the knowledge gained by participating in these projects. With these multiplier events, we hoped that the teachings and the resources of this project will be showcased to a wider audience; thus, enabling the advancement of Ethical AI teachings across universities in the EU. Dissemination is an important part of any project and with Trustworthy AI Multiplier Event, we hope that the attendees will be inclined to adopt the developed resources as part of their teaching methods.

Photo by Conny Schneider on Unsplash

Co-funded by the Erasmus+ Programme of the European Union

Photo by NUIM

Partners experience in running the hackathon

Maynooth University, Ireland organized the Ethical AI Hackathon as part of the Trustworthy AI Project over the period of 15th to 17th June, inviting third-level students from all disciplines across Ireland. With participation of over 30 students, the hackathon managed to successfully introduce the students to the ethical and value-based components of Artificial Intelligence and use their learning towards a fascinating use case in the healthcare domain. Students formed a team of three or four and were aided by a mentor. The project involved faculty members mentoring the students during the course of their participation.

The three-day virtual hackathon was hosted over devpost and aimed to introduce the participants to the resources including the introductory videos, the card deck and the 7-step exercise to assess an ethical dilemma in AI on day 1. Day 2 began with an introduction to the use case, CHAPE ((http://cohealth.ivi.ie/chape/), a health empowerment platform that lets users interact in human-readable language to enter their health data and symptoms to generate relevant information from credible sources. The last two days of the hackathon were dedicated to building an ethical framework for CHAPE that could address all the seven requirements of Trustworthy AI. The virtual hackathon encouraged collaboration between students via Padlet, a collaborative tool that supports all major filetypes and lets one embed media from all major sources.

Four submissions were shortlisted for the final judgement wherein Dr. Marco Alfano, Dr Markus Helfert and Dr. Zohreh Pourzolfaghar from Innovation Value Institute, the research hub for digital innovation in Maynooth University, selected the winning team which received One4All vouchers along with recognition and felicitation at the IVI Summit on 30 June 2022. The Ethical AI Hackathon proved to be a successful venture to introduce the students to Trustworthy AI and help them apply their learnings to a live project.

Featured articles

Photo by Tyler Lastovich on Unsplash

Trustworthy AI

Co-funded by the Erasmus+ Programme of the European Union

Teaching Students About the Challenges of Artificial Intelligence

When we talk about Artificial Intelligence, it is important to understand that its pillars are three distinct academic disciplines: psychology (cognitive modelling), philosophy (philosophy of mind), and computer science. Other fields that also contribute to the field of Artificial Intelligence are linguistics, mathematics, and logic. Joining these very different areas is far from easy. That is why teaching Artificial Intelligence in Higher Education is a complex assignment, but our Erasmus+ project is here to support this cause.

As we know, the aim of Artificial Intelligence is primarily to discover the processes, systems, and principles that enable intelligent behaviour. The science behind Artificial Intelligence was once just wishful thinking, but today we see intelligent behaviour in machines all around us: drones, self-driving cars, manufacturing robots, disease mapping, chatbots, etc. Artificial Intelligence is already a part of our lives, and it is on its way to become a regular part of the education, too. However, if we aim to include AI in higher education, there are a few aspects of the field that are key to be included in the curriculum.

Students should be able to understand the opportunities, as well as the challenges posed by AI. The challenges are increasingly posed by the adoption of Artificial Intelligence in our everyday life. Democratic societies cannot have AI included in education without careful consideration of the ethical issues ingrained in the very essence of AI. Companies which design AI programmes are not

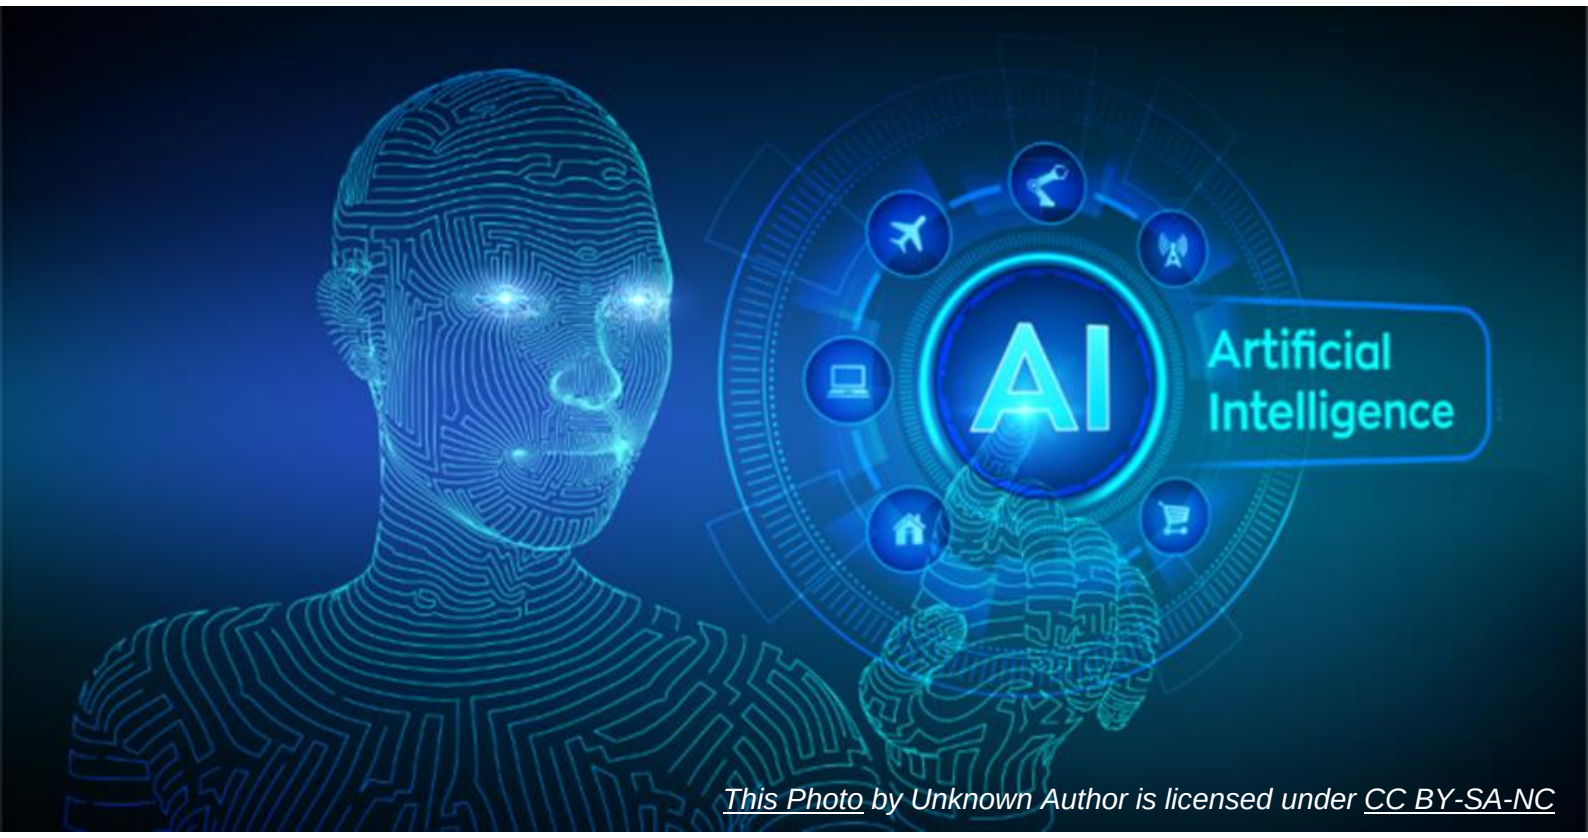

Co-funded by the Erasmus+ Programme of the European Union

the only ones that are facing these issues. Organisations which use the AI data face ethical issues, as well, and are often unfamiliar with the proper way of addressing them.

The course of AI should ideally discuss not only the science behind Artificial Intelligence, but also global governance and corporate social responsibility, since a company using AI has not only a moral, but also a legal obligation to responsibly deal with the data. Therefore, including the legal framework that deals with Artificial Intelligence in business is also an essential part of the AI education.

Artificial Intelligence poses another challenge, which is seen in the inequality of resources, opportunities and power in the workplace. Students studying Artificial Intelligence need to be aware of this issue and become familiar with the tools that can be used to overcome it.

If students learn about Artificial Intelligence, their learning experience should not be limited to just adopting the information from the course, but they must also be equipped with critical thinking skills in order to able to analyse and discuss these ethical issues.

It is absolutely clear that education plays and will continue to play a critical role in the evolution of Artificial Intelligence and its integration in the business world. As the society in which we live in is rapidly expanding its use of Artificial Intelligence, the academic world must keep up. There is an evident gap between the two currently and including Artificial Intelligence in the higher education curriculum is an imperative for the future.

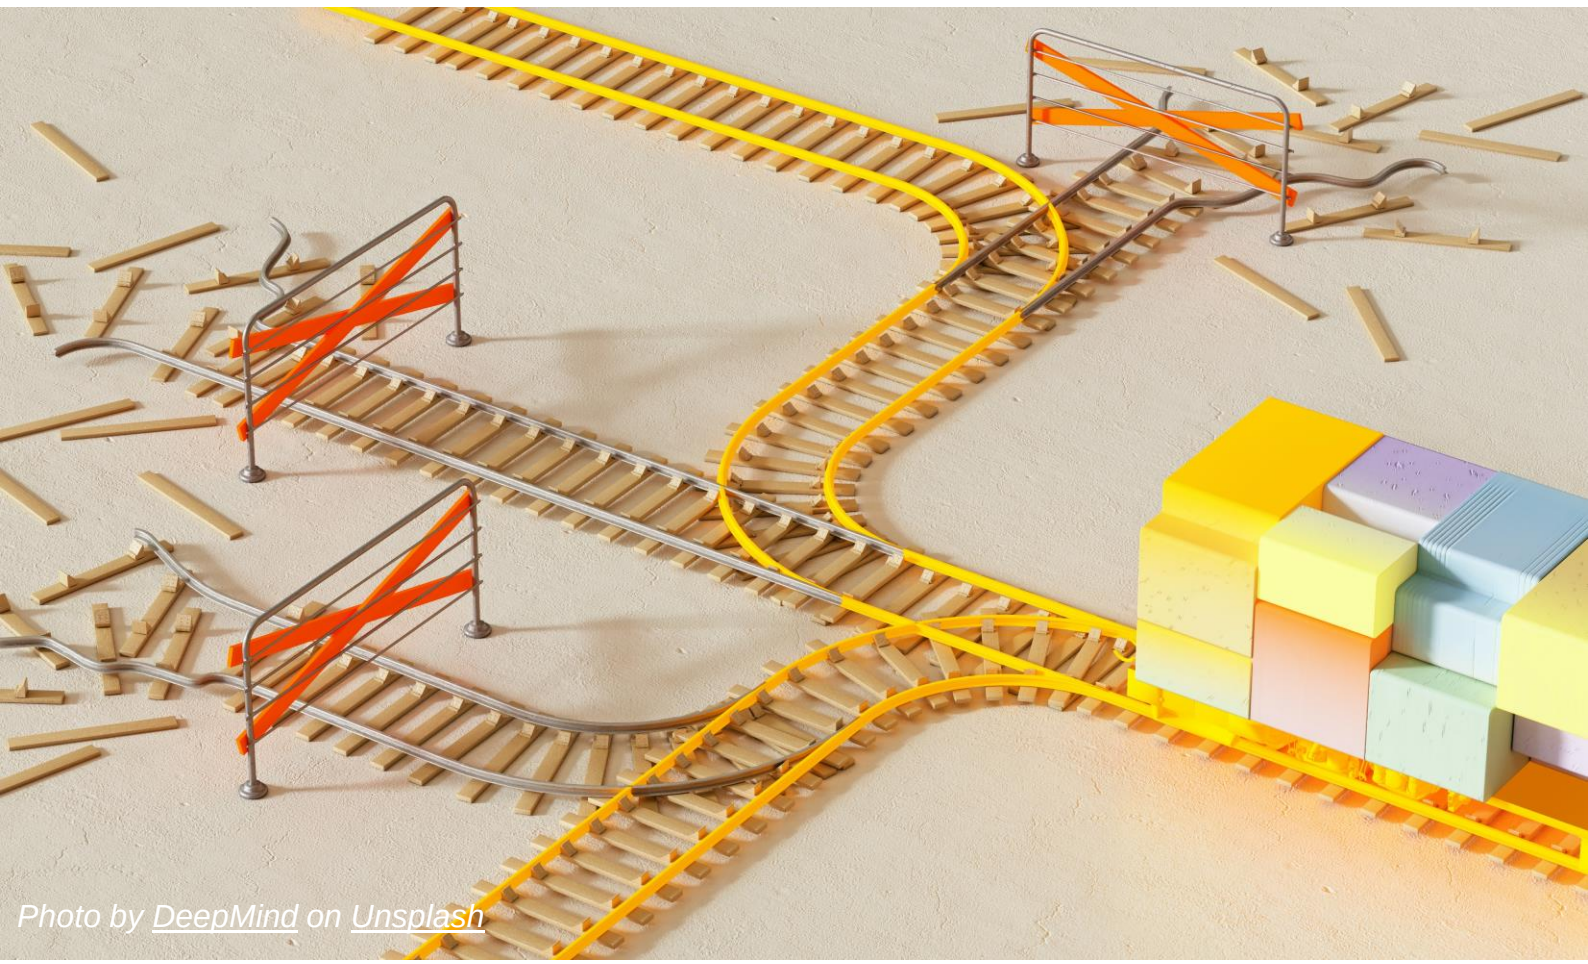

Photo by DeepMind on Unsplash

Trustworthy AI beyond the project

What is the potential of our project resources for the future?

It's been two years since we started with this project and the only thing that has not changed is that time flies. Everything else has changed somehow due to the project. For instance, the consortium. We started slowly at the end of 2020, unsure of where this was going to lead us. We had strong foundations, with some very experienced partners and others that were experts in their respective fields, but the topic (Trustworthy AI) and the environment (HEI) inspired some respect. The partnership started to work since day one; it is of course true that there have been some highs and some lows, but overall, the collaboration has been outstanding – and I'm completely sure that some future long-lasting bonds have been formed, which will give birth to new proposals and projects in the future. All institutions have been impacted on one way or the other by the project and this is visible in its results.

The objective of the project was clear since the start: our plan was to adapt the HLEG AI requirements and the Assessment List to be able to teach the required competences in HE institutions. The road, however, was new. With the first IO, the Framework, led by UMU, we noticed how there was much work to do and the long list of potential gaps and risks for the project. Interviewing experts from all regions helped, but they also confirmed our fears – it was hard to think about changing the curricula or the politics at this stage, so the best course of action at this point would be to empower teachers of all disciplines to include Trustworthy AI in their own teaching with flexible and adaptable materials. The Framework has had hundreds of downloads already, which shows its relevance and how it was needed.

IO2, the Open Educational Resources, also started with a warning: a review of the existing materials resulted in a very short list of available resources. Therefore, we had to start from scratch, so ALLAI and EUEI partnered to bring us three amazing teaching tools in the form of a set of videos, a deck of cards and the 7-step method, which are not only freely distributed but they can also be adapted by any teacher to their topic or discipline of choice. The feedback in the user tests was very positive, so the impact expected in all HEI institutions is huge. As the project comes to an end in August, it is key for us to ensure its sustainability, to increase the expected impact after the end of the project.

Finally, the third output, the Guide to Conducting Ethical AI Hackathons, could be seen as a fourth – but larger – resource, that includes not only the suggested way to use this methodology or tool to introduce Ethics, Trustworthiness and AI to any public, but also contains and makes use of the previous resources to multiply their impact. The Guide was used by three of the parners to conduct their respective pilot Hackathons and the response was, again, inspiring.

As said, the project ends in August, but it is not the end for its results. The consortium will keep disseminating its contents to HE institutions, while building on top of the existing collaboration networks to multiply its impact. The OERs will be kept alive, and help will be provided to those teachers who wish to adapt the provided resources to their particular needs. Each partner, in turn, will become a national point of contact for any questions related to the project or the topic, whenever required. And the website will be kept operational at least for the next few years. We are also actively pursuing further opportunities to extend or expand the project with new proposals, projects and partnerships with the aim of helping to empower our citizens to build a Trustworthy AI made in Europe.

Co-funded by the Erasmus+ Programme of the European Union

Trustworthy AI

CONTACT US

Marçal Mora Cantallops
Project Coordinator, UAH
marcal.mora@uah.es

Mario Ceccarelli
Dissemination & Outreach, UIIN
ceccarelli@uiin.org

Trustworthy AI project

Trustworthy AI project

trustworthyaiproject.eu