

Should We, As Researchers, Care about the “Magnifica Humanitas”?

Matteo Cardellini¹

¹Università degli Studi di Genova, Via all'Opera Pia 13a, Genova, 16145, Italy

Abstract

Artificial Intelligence (AI) is rapidly changing our world and steering it through an opaque direction. Since the introduction of Large Language Models (LLMs) to the public, massive investments have been made into the sector. This surge in interest in the field is also noticeable in conferences, where the number of papers submitted is ever-growing. While these conferences present cutting-edge programs, with very technical and novel papers, the author's position is that they fail to provide researchers with tools to develop a moral conscience, an ethical perspective, and a societal framework to help them navigate, as first-hand actors, through the theories and applications they build every day. On the other hand, the discussion on morality, ethics and societal impact is done far away from technical conferences: Pope Leo XIV, far from being a technical expert on the subject, wrote the encyclical letter “Magnifica Humanitas”, focusing brilliantly on how to safeguard humanity and the integral human development in the age of AI. This discussion paper argues in favour of dialoguing with philosophers, ethicists, and religious leaders, bringing insights from the human sciences, hearing stories about how society is impacted by AI, with the aim of helping us to build a moral compass for our research activities.

Keywords

Ethics of Artificial Intelligence, Society, Magnifica Humanitas

Context. Artificial Intelligence (AI) has recently gained widespread popularity among non-experts, especially after the public release, in late 2021, of ChatGPT, i.e., one of many Large Language Models (LLMs) now available. This technology is rapidly changing our world and directing it through an opaque direction. Indeed, while, in the past, nations and public institutions were driving and introducing new technologies, nowadays, they have been replaced by private institutions and companies, which have far superior economic power and financial resources than many governments. While one could argue that most research projects in the past were indeed public, but under the defence budget, and thus, under secrecy (see, e.g., ARPA for the invention of the internet), it is true that their purposes and applications were (at least in theory) to be governed by elected officials, mandated democratically by the people. Now instead, major AI research and applications are conducted by private companies, where projects are under NDAs and trademarks, and where managers have obligations only towards shareholders and not to the public, and the direction imposed is not always directed towards the common good.

This paradigm shift is also affecting universities and research institutions. Especially with LLM research – which is now occupying a large portion of the research output in many conferences – where training the models requires expensive technical infrastructure, the possibility of participating in the research is limited to wealthy research groups that can afford such infrastructures. Moreover, research on the frontier models and how they operate is mainly done by calling pre-existing (and mostly private) models through APIs, since training one from scratch can only be achieved with gigantic infrastructure, outside the reach of any research institution.

In this rift between public and private, the world is currently debating important moral questions regarding AI and its societal impacts: (i) the use of AI in warfare, (ii) the role of AI in children's education and psychology, (iii) the loss of jobs being replaced by AI, (iv) the gigantic use of natural resources in data centres training AI models, (v) the fairness of these models in risk-sensitive fields like medicine,

AIxLA 2026 Discussion Papers, 24th International Conference of the Italian Association for Artificial Intelligence, Perugia, Italy, 2026

✉ matteo.cardellini@unige.it (M. Cardellini)

🌐 <https://matteocardellini.it/> (M. Cardellini)

🆔 0000-0003-3788-9475 (M. Cardellini)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

justice and economy, (vi) the gap between rich and poorer countries growing more if the latter has no access to AI technologies, and many others among these.

The Role of Conferences. Conferences¹ are events where researchers meet to present their work, get feedback, brainstorm ideas with others and build networks. Conference submissions are steadily increasing (e.g., AAAI received more than 23.6k submissions in 2026, almost twice as many as the 12.9k submissions it got in 2025). Conferences' programs present a plethora of technical papers of extreme quality, and keynote speakers (highly respected researchers in the AI field or in neighbouring ones) provide a technical compass towards where to point one's research.

In a world where international cooperation between governments is at its minimum, the United Nations feels powerless in front of the global problems, and multilateralism is no longer considered an option, international conferences provide a rare opportunity for dialoguing across borders, to get informed about other people's opinions, and to take a stand for what we believe to be right and important. It is my position that these conferences would be the perfect place to provide researchers with tools to develop a moral conscience, an ethical perspective, and a societal framework to help them navigate, as first-hand actors, through the theories and applications they build every day.

“Publish or Perish” and Feeling Powerless. As a young researcher myself, I feel sometimes crushed by the mantra “Publish or Perish”, which brings a young researcher to only think about how and where to publish, and to fill one's résumé as fast as possible, with the beauty of research, and the joy of the sense of discovery being left apart. With this mindset, interest in ethics and morality, together with understanding the societal framework where one's work applies, seems counterproductive, a waste of time, and thus is rarely explored, since it would shift the focus from one's main research. Moreover, one can easily feel powerless: a small cog in a much larger mechanism and without any actual leverage to change the state of things. For this reason, most people remain on the sidelines, waiting, observing from afar, hoping that everything will sort itself out.

Talking about ethics in conferences could help in solving both aforementioned problems: (i) a researcher who attends a conference is already there to present or attend technical sessions, and thus would not seem as a waste of precious time, and (ii) being together, engaging in dialogue and sharing each other's opinions is the first step towards building something and leaving the loop of powerlessness.

Most importantly, we have to care. An Italian educator, Don Lorenzo Milani (1923-1967), in '54 opened a school for children of farmers in Barbiana, in the mountains of Florence, and was the first to realize (at least in Italy) that school was the only way for poor people to rise from their social condition. In the classroom, which is still standing, it is written, in capital letters, the English phrase «I care». The phrase was a response to the phrase «I don't care» (in Italian, «*Me ne frego*») which was typical in the fascist era, to denote the disinterest towards other people's problems. As young researchers, we feel the weight of an environment telling us to «don't care» and to «think of yourself first». Don Milani was convinced that feeling responsible, caring, implies the ability to go beyond oneself to take on the problems of others. As engineers and scientists, we think that a problem can be solved by *doing*. Since ethical and moral problems seem bigger than us, we can tackle the problem first by *caring*. Caring, together, will resolve the sense of powerlessness and guide us towards a solution for the problem.

Magnifica Humanitas. Ethical and moral considerations concerning AI take place, however, far away from technical conferences. Robert Francis Prevost's first decision, after being elected pope in May 2025, was to choose the papal name of Leo XIV. His predecessor, Leo XIII (1878-1903), was elected pope during the First Industrial Revolution and was concerned with the human dignity of workers. Prevost felt that AI was bringing another industrial revolution, and as pope, he had to deal with human dignity in the AI revolution [1]. His first Encyclical Letter, i.e., the letters in which popes present their considerations on what they believe to be important, is titled “Magnifica Humanitas - On Safeguarding

¹I periodically attend international conferences like AAAI, IJCAI, KR and ICAPS

the Human Person in the Time of AI ” [2]. These choices follow a long pattern of discourses by his predecessor, Pope Francis, even before AI became mainstream [3].

The fact that popes, who are clearly not technical experts on the subject, felt the urgency to speak about AI should at least raise interest, and show that the problem must be taken seriously. Clearly, it is the job of a pope to guide the moral standing of Christians. It is our job, instead, as experts in the field, to build and advance these technologies. It seems thus we are at a crossroads: here, technology and there, humanism. Should these two roads be totally separated, or should we join forces and teach one another? This position paper argues for the latter. In the rest of this paper, I’ll report what I learned from reading the encyclical letter, and what, I believe, could be beneficial for the AI community.

Two Ways of Building. In the encyclical, the main focus is not on stopping building AI technologies, but on *how* we build them. The pope, with a biblical reference, presents two ways of building ([2], 7-10): (i) the construction of the tower of Babel, which is done fast, without any care, to get power and personal gain, sacrificing dignity, or (ii) the reconstruction of the city of Jerusalem, which is done together, giving each person a role in the reconstruction, comforting scared people, coordinating the works and engaging in dialogue with opposing views. Thus, not a development done by the few, the most wealthy companies in the most wealthy countries, but a development where responsibility is shared, between companies and research institutions, between richer and poorer countries. Where building is done together, transparently, and not for building wealth, but to improve human dignity. The difference is between separate individuals building separately, or of a community that builds together.

A Human Development. Thus, the “Magnifica Humanitas”, i.e., the magnificent humanity, which includes researchers, but also private companies, schools and politicians, is called to decide whether to reconstruct a city or to build Babel, to have progress that serves the people or that bends them to the logic of power ([2], 129). The encyclical calls for the development of AI to be *human* ([2], 83), i.e., «when it places people at the centre instead of the accumulation of wealth, and when it concerns peoples as well as individuals». Computer scientists, engineers, and mathematicians are very good at reducing complex problems into formulas, data structures and diagrams. While this is helpful for solving the problem, one must never forget the humanity that lies beyond the numbers. Feasibility of a solution must also consider the human factor, and not only the physical constraints of the problem. When multiple solutions are possible, the optimal one must not be computed with metrics that only consider profit, but that also considers the well-being of a person.

The pope then addresses a special appeal to us ([2], 111): «Developers, therefore, bear a particular ethical and spiritual responsibility, for every design choice reflects a vision of humanity. Just as the creator of an artistic or literary work must consider the values it conveys, so developers are called to embed values in their projects with due seriousness: with transparency, responsibility toward affected communities and careful attention to ensuring that what is being cultivated is a genuine good». The title of this paper is thus also an appeal to care about the magnificent humanity in our works, to never forget the humanity behind the numbers. This is only possible if we care (as Don Milani said), if we allow ourselves to be moved by others, if we give faces to the numbers, if we hear their stories, if we recognize their wounds.

Climate. Climate change is one of the most pressing issues facing our world today. Over the last 50 years, global temperatures have risen significantly, with the Earth’s average surface temperature increasing by about 0.9 °C to 1.2 °C since the late 19th century. AI is unfortunately contributing to this phenomenon. Training LLMs requires a lot of power to run and water to cool the hardware. Big companies are training LLMs on larger and larger datasets, increasing their parameters exponentially every few months. Since most capital investments are done in the field of AI, for big companies, the amount of money to be spent on new data centres and infrastructure appears to be virtually unlimited. Companies producing computer parts, like GPUs, sell astronomically large quantities of components, which require huge amounts of rare metals which are extracted from the earth.

As researchers, we have to focus on searching for more viable and less environmentally impactful ways for LLMs to be trained. We must stop using LLMs technologies as a Swiss Army knife for every problem, and reflect on the most appropriate and sustainable tools to solve the problem. We have to dispute the idea that throwing more computational power, more tokens and more data into a problem is the right way to solve it.

Warfare. Conflicts and wars are at the highest level² since the end of the Second World War (WWII). A “culture of power” is consolidating ([2], 188-208), where war is rehabilitated as a tool of international politics while the ethical principles, born at the end of WWII, put in place after recognizing the terrible consequences of war, are eroded. Humanity is slipping in the violent culture of war, getting accustomed to it, where peace is no longer considered a goal to achieve but a precarious interval between conflicts. In this context, we are getting accustomed to hearing about AI being used for autonomous weaponry, for military logistics and for decision-making. As mentioned in the previous paragraph, in warfare applications, human life is reduced to data and to a metric in a cost function. The pope proposes real points of discernment in the use of these applications ([2], 199). The first has to do with personal responsibility: when the decision to strike is automatic, there is the risk of not feeling personally responsible, and thus bypassing the ethical and moral question needed for such decisions: «Any technology that facilitates attacks without seeing the face of human beings lowers the moral threshold of conflict». The second point has to do with time: AI tools are mainly used as a way to speed up decision-making and to delegate tedious tasks. Thus AI tends to compress the time to decide. However, in war, irreversible decisions cannot be judged only by criteria of speed and efficiency. There is the need for time to discern and evaluate other solutions, where the use of force is mitigated and life preserved. The third criterion has to do with the protection of civilians. Target selection and the use of force must not confuse combatants and non-combatants, nor ignore the impact on defenseless populations.

The big hype in AI of the last years, mostly concerning LLMs, is making the public (and most importantly, politicians) think that LLMs can solve any problem they are tasked to solve. As researchers who know the theoretical aspects of AI systems, we have a profound knowledge of their merits but, most of all, of their limitations. We know that LLMs sometimes hallucinate, provide wrong results and have a limited ability to reason. In warfare applications, these limitations could mean loss of human life that could have been prevented. As researchers, thus, we have the duty to keep the conversation about AI based on the facts, avoid increasing the hype on these technologies and inform the public of their limitations.

We can all do our part. As stated some paragraphs above, we certainly feel powerless in front of all the aforementioned problems. We think that the problems are too big and we are too small. However, no one is without responsibility, even in our small lives. «We all have our own areas for action, and it is precisely there — and nowhere else — that we must choose whether to fuel the mentality of force (even if only through indifference, cynicism, lies or hatred), or to preserve the mindset of peace (with truth, moderation, closeness, and care)» ([2], 212). There is no need for grand and spectacular gestures, but for small but faithful and convinced contributions. For my part, this paper was one of these small contributions. I am not arguing that AI conferences should become ethics conferences, and technical excellence should remain the primary mission of AI conferences; however, technical excellence alone is no longer sufficient. Conferences define research agendas not only through accepted papers, but also through the conversations they choose to host. I invite program chairs to include more discussion papers like this in their conferences and to include, in their programs, space for dialogue, where everybody can express their opinion, worries, and hopes. I also encourage them to include invited talks by researchers from the human sciences: psychologists, ethicists, political scientists, philosophers, educators and religious leaders. I invite conferences to act as a multilateral organization without borders, where the worldwide research community is represented. I appeal to researchers to care about the magnificent humanity, so that, by caring, they bring these insights into their work and applications.

²<https://ucdp.uu.se/encyclopedia>

Declaration on Generative AI

During the preparation of this work, the author used ChatGPT5.5 and LanguageTool to: (i) Improve writing style, and (ii) Grammar and spelling check. After using these tools, the author reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] Pope Leo XIV, Address of the holy father to the college of cardinals, 2025. URL: <https://www.vatican.va/content/leo-xiv/en/speeches/2025/may/documents/20250510-collegio-cardinalizio.html>.
- [2] Pope Leo XIV, Encyclical letter *Magnifica Humanitas*: On safeguarding the human person in the time of artificial intelligence, 2026. URL: <https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html>.
- [3] Pope Francis, Speech at the meeting on promoting the digital dignity of minors, 2019. URL: <https://press.vatican.va/content/salastampa/it/bollettino/pubblico/2019/11/14/0868/01817.html#in>.