

THE FUTURE OF CATHOLIC UNIVERSITIES IN THE AGE OF AI:
A ROADMAP OF CENTRAL ISSUES

SACRU Scientific Colloquium
Università Cattolica, Milan, Italy,
July 2023

Micah Lott, Boston College
William Hasselberger, Universidade Catolica Portuguesa
Francisco Nunes, Universidade Catolica Portuguesa
Federico Manzi, Università Cattolica del Sacro Cuore
Marco Carlo Passarotti, Università Cattolica del Sacro Cuore

Given AI's increasing power and transformative potential, the SACRU conference was an occasion for examining not just the benefits and risks of a specific technology but the basic purposes and aspirations of a Catholic university in a technologically advanced society. At its best, what might a Catholic university be? What are the distinctive goods that Catholic universities can achieve and that give orientation and coherence to their activities? What are the essential and defining roles within a university, and what virtues and skills are required for persons to succeed in those roles as members of a university community? In answering such questions, this paper develops what might be called a normative vision for Catholic universities in a technological age. Because this was a meeting of Catholic research universities, the normative visions of the university can be divided into two broad areas: teaching and research. Because this was a meeting of Catholic research universities, the normative visions speak to the distinctively Catholic approach to higher education, including a conception of spirituality, the dignity of human persons, and the university's proper role in furthering the common good of society.

The position paper adopts a realistic approach that addresses two challenges: the educational and the anthropological ones. It is divided into sections identifying the most pressing issues for the Scientific Colloquium. The issues are organized under three headings, based on the central tasks of the Catholic universities as understood by the authors: teaching, research, and service to society. The point is not to provide a rigid taxonomy but to give some organization to this complex technological revolution. In each section, the authors try to be aware of both the benefits and risks of AI in universities. Questions about possible connections between AI and the normative visions of the Catholic university are also included.

A common Vocabulary: A definition of AI

In the context of research and reflection on Artificial Intelligence (AI), the necessity of establishing a shared definition is a foundational starting point. This need arises from the widespread use of the term 'AI', which can have a semantic spectrum broad enough to introduce ambiguity. Moreover, in the AI context, there is an even more pressing need for a common terminology, given the high multidisciplinary nature of the field and the convergence of different communities. From lawyers to philosophers, theologians to linguists, engineers to mathematicians, physicists to psychologists, and economists to sociologists, AI attracts experts from various disciplines.

During the recent SACRU conference held from July 13-14 in Milan at the Università Cattolica del Sacro Cuore, Xavier Vilasís (Universitat Ramon Llull) proposed an AI definition that can be considered an apt summary of AI's fundamental characteristics: "A set of mathematical techniques that perform tasks we associate with human beings". This definition succinctly captures three fundamental elements: computation (the "mathematical" aspect), operation ("tasks"), and the human role ("human beings").¹ However, the very use of the concept of "intelligence" in a technological context can be potentially misleading if we take human intelligence as our point of reference. AI systems operate quite differently since they typically perform prediction or classification tasks by applying statistical techniques to large datasets to identify recurring patterns. This computational operation can discern data connections that are hidden or invisible to the human eye (a concept known as the "black box"), with significant consequences concerning the im-

¹It is important to distinguish AI in this sense usually called 'algorithmic' from the so called Artificial General Intelligence. This is AI that not only performs tasks associated with human beings, but that also rises to human level intelligence (e.g., by passing a sophisticated form of the Turing test or perhaps by possessing traits that go beyond). Whether such AI will be produced in the future is open to debate. At

pact of AI's application in societies when delegating decision-making tasks based on knowledge to AI (Bengio et al., 2023)². Nonetheless, the provided definition seems to lack a deeper understanding of AI's impact on humans. This impact dimension was well outlined by the AI Act³:

"...software that is developed with [specific] techniques and approaches [listed in Annex 1] and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with".⁴

Assessing this impact on humans and society implies a realistic approach to interpreting AI's role. This is the approach pursued by SACRU in addressing the advent of AI in human life.

I Introduction

AI and a Normative Vision for the Catholic Research University

This conference of the Strategic Alliance of Catholic Research Universities (SACRU) explored the special possibilities and challenges for Catholic universities that arise from Artificial Intelligence (AI). The history of AI since the 1950s has been a history of alternating periods of enthusiasm, lavish funding, and high expectations of immanent breakthroughs in machine intelligence ("AI springs"), followed by periods of dashed expectations, unfulfilled promises, and funding cuts ("AI winters"). But things seem genuinely different this time. In the last ten years, there have been major advances in AI image recognition, complex game-playing, and, most recently, Large Language Models (LLMs) and other forms of "generative" AI. Hundreds of millions of users now rely on AI-powered systems, including algorithmic content feeds on social media and

any rate, AI of the type defined above already carries the potential for significant benefit and harm, and so it is this form of AI that we focus on.

²<https://managing-ai-risks.com/>

³<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>

⁴ OECD, Recommendation of the Council on Artificial Intelligence, 2019.

the Chat-GPT chatbot. Next-generation AI systems are demonstrating impressive and sometimes unanticipated powers. A recent study by a team of researchers at Microsoft claims that the GPT-4 language model (the more advanced version of the language model underpinning the Chat-GPT system) is showing “sparks” of truly “general” and “human-level” intelligence—the long-standing dream of AI research—and will thus transform work and labor-markets, and perhaps social life more broadly. Across industries and sectors of society, more and more cognitive tasks and decisions that we were previously the responsibility of human beings are being automated by AI.

AI breakthroughs are raising serious concerns about the value and future of higher education. A widely-shared article in *The Atlantic* boldly (if sorrowfully) pronounced that “The College Essay is Dead.” Others have argued that next-generation AI will render many traditional academic disciplines obsolete. Still others have argued that the ethical challenges from AI, including dangers of mass-scale automated misinformation, show that strengthening university education is crucial to safeguarding democratic societies.

Given the increasing power and transformative potential of AI, the SACRU conference was an occasion for examining not just the benefits and risks of a specific technology, but the basic purposes and aspirations of a Catholic university in a technologically advanced society. At its best, what might a Catholic university be? What are the distinctive goods that Catholic universities can achieve, and that give orientation and coherence to their activities? What are the essential and defining roles within a university, and what virtues and skills are required for persons to succeed in those roles as members of a university community? In answering such questions, we develop what we might call a *normative vision for Catholic universities in a technological age*.

Without such a vision, our discussions of AI are liable to be misguided in various ways: overlooking important questions, overly focused on technical and regulatory details, too pessimistic or too optimistic, etc. Hence, the participants of the conference have been encouraged to bring their own normative visions of a Catholic university to the table, and be ready to question, re-examine, and refine those visions in conversation with others.

Because this was a meeting of Catholic **research** universities, our normative visions of the university can be divided into two broad areas: teaching and research. Because this was a meeting of **Catholic** research universities, our normative visions will need to speak to the distinctively Catholic approach to higher education, including a conception of spirituality, of the dignity of human persons, and of the university’s proper role in furthering the common good of society.

With this in mind, here is an initial, and very general, proposal for a normative vision for Catholic research universities. Universities exist for the good of those who participate in them, and for the larger societies of which they are a part, by pursuing the distinctive goods of **understanding** (knowledge, insight, wisdom). Within the realm of *teaching*, this especially includes passing on knowledge to students, helping them to acquire a body of knowledge for themselves, as well as the skills and intellectual virtues necessary for meaningful, important, deep understanding of themselves and the world (both the human, cultural world, and the natural world). Within the realm of *research*, this especially includes the grasp of specialized knowledge in various disciplines, and the generation of new knowledge, as well as new means for acquiring knowledge. Within the domain of *service* to the community, this especially includes disseminating understanding, and fostering public discussion of both perennial ideas and contemporary challenges, with an eye to respecting the dignity of all persons and fostering the common good.

Now, how might emerging AI technologies impact—for good or ill—the mission and purpose of Catholic research universities?

A Realistic Approach. Two challenges.

In thinking about the impact of AI, we need an approach that is *realistic*. This has two dimensions. First, we need to think about both short-term and long-term impacts. AI has already begun to shape both teaching and research, from automated plagiarism detection software using machine learning, to student papers written with “assistance” from Chat-GPT, and academic research articles that have used—and listed—generative AI systems as a “co-author.” Recently, even the writing process of proposals for research funding has been performed with the help of AI-based systems (Parilla, 2023)⁵. Even more significant impacts are likely close at hand. Thus, we need to address pressing, contemporary concerns about automated plagiarism, student evaluation, author’s rights, erroneous outputs of generative AI systems (so called “digital hallucinations”) in research papers, the “death” of the essay, and so forth.

At the same time, we should also consider the more sweeping impacts of AI that might be farther in the future, including profound social changes to work, leisure, and political and community life. These might be harder to predict and to understand, but nevertheless worth pondering, because of their potentially drastic effects. We need to be realistic about the changes that are already occurring, and not focus only on futuristic, science-fiction scenarios. At the same time, we need to reflect on those even greater changes that might be on the horizon, recognizing that, in some cases, what sounds fanciful today will be reality in the future.

Second, realism requires us to think about both the potential benefits and potential harms of AI for Catholic universities. It is unwarranted

to assume that AI will have only, or even primarily, positive impacts on teaching and research. At a time when leading researchers in AI like Geoffrey Hinton, Yoshua Bengio, and others, warn of potentially catastrophic social effects of next-generation AI, we should not be naïve cheerleaders for an AI future. At the same time, we should not assume that AI will have only negative impacts on our universities. Realism about AI means having curiosity and openness to its potential benefits, rather than a knee-jerk pessimism or doom-and-gloom defeatism.

To fully grasp the impact of AI and avoid passive consumption of it, it is crucial to understand what AI is. AI is not just a matter of technology; it has the power to influence the decisions people make based on the outcomes produced by AI systems. In order to better understand those outcomes, and to maximise their benefit, we must know the procedures by which they are generated before they manifest in our reality. This will enable us to make informed decisions about how to use AI, as well as how to respond to the use of AI more generally. The conference has highlighted at least two significant challenges that contemporary society will need to address, challenges to which Catholic Universities can provide a fundamental contribution, given their mission. The first represents an educational challenge, and the second is an anthropological challenge. Below, we will delve into what these two challenges mean for Catholic Universities.

Educational Challenge: The educational challenge addresses the need to teach students, scholars, professors, and society at large how AI works before incorporating AI components into educational processes and daily life in general. This is to avoid the risk of a “ChatGPT said!” scenario leading to the spread of fake news and misinformation. Education about AI

⁵ Parrilla, J. M. (2023). ChatGPT use shows that the grant-application system is broken. *Nature*.

should precede education with AI. This education about AI should be included in all school programs, starting from primary schools, to bridge the gap between highly specialized AI professionals (who are active consumers and developers of AI) and society (which is a passive consumers of AI). Students should be trained to exercise "digital discernment", promoting a culture of reflection and critical thinking in an environment heavily manipulated from an informational standpoint. This poses one of the great challenges to democracies and is also a challenge for universities, the cradle of democratic thinking. Universities should support basic education about AI by providing appropriate skills and tools. Concerning higher education, finding a balance between excessive caution and reckless innovation is necessary. We should consider that advancements in AI technology are transforming the job market, changing the skills employers value in their employees. This implies that universities should start considering how to align their educational programs with the needs of the job market to support graduates and equip them with the necessary training and knowledge for today's technological landscape. To this end, the breaking down of disciplinary boundaries and the introduction of highly interdisciplinary bachelor's and master's degrees are crucial, leading to the next challenge.

Anthropological Challenge: The anthropological challenge involves deep reflection on human nature and how AI can influence it. This challenge lies at the heart of Catholic universities, which traditionally place particular emphasis on human dignity, understanding the role of humans in the world, and the promotion of

ethical values. If we do not successfully address this challenge, we risk losing sight of the essence of being human. This can lead to several negative consequences, including detrimental effects on thought processes, as excessive reliance on AI and the unaware use of these automated devices could erode our ability to concentrate, think critically, and tackle complex problems, resulting in a kind of "atrophy" of human mental faculties. On the other hand, if we successfully tackle this challenge, there are significant advantages to be gained:

1. *Self-Awareness:* A profound understanding of how AI works and how it interacts with our daily lives can contribute to greater self-awareness. This awareness can aid in better managing AI consumption for personal well-being and societal benefit.
2. *Enhanced Spiritual Education:* The automation of language, as suggested by Roberto Busa (1958)⁶, requires an advancement and deepening of human spiritual education. Having a machine perform functions typically associated with humans, such as language and, more broadly, intellectual processes, first and foremost requires humans to better understand these functions, which essentially means knowing oneself on a deeper level than ever demanded by humanism.
3. *Emphasis on Humanity:* This challenge encourages Catholic universities to emphasize humanity. Despite and thanks to technological advancements, humans remain at the core of the educational process and ethical decisions.

⁶ This leads me to conclude that the automation of written language awaits some technical development, but it also expects much more from the spiritual industriousness of mankind. The machine warns us that we are not humanistic enough and, although we speak, we are not able to explain how we speak. It is the despised machine that repeats to us the invitation "know thyself still more profoundly, scientifically and humanistically: study your speech". The automation of written language thus promises an increase in spiritual education. (English translation

taken from Nyhan and Passarotti, 2019, pp. 68; originally published in Roberto Busa 1958).

Julianne Nyhan and Marco Carlo Passarotti (Eds), One origin of digital humanities. Fr. Roberto Busa in his own words. Springer, Cham. Germany. 2019.

Busa, R. 1958. I principali problemi dell'automazione del linguaggio scritto. In Atti della VI Sessione delle Giornate della Scienza – Convegno Int. sui Problemi dell'Automatismo, Milano, 8–13 Aprile 1956, vol. I. pp. 831–841. Roma: Consiglio Nazionale delle Ricerche.

4. *Ethics and Inclusivity*: Algorithmic ethics that respect plurality and promote inclusivity become crucial. AI can be a means to improve human interaction and give voice to diverse perspectives and experiences.

In summary, the anthropological challenge in the age of AI calls for deep reflection on human nature, ethical AI consumption, and the importance of valuing humanity in an increasingly technological world. Catholic universities can play a fundamental role in promoting these ideas and in shaping individuals who are aware, ethical, and capable of addressing the challenges of the AI era.

Organizing the Issues

The sections that follow identify what we take to be the most pressing issues for this conference. We have organized these issues under three headings, based on the central tasks of the Catholic universities as we understand them: teaching, research, and service to society. In some cases, the issues will clearly belong to multiple headings, so the assignment to one is somewhat arbitrary. The point is not to provide a rigid taxonomy, but to give some initial organization.

In each section, we have tried to be aware of both potential benefits and potential harms of AI in universities. We have also included some questions to bring out possible connections between AI and our normative visions for the Catholic university.

II Teaching, Learning, and Student Formation

First, contemporary AI systems, and especially generative AI tools, depend upon the massive collection of data from users and the wider internet, often without the consent of, or financial compensation for, the producers of the underlying training data. Serious ethical concerns

exist regarding **data protection and use** (especially commercial use), and there have been instances of professors' and students' private information being violated or used for purposes of which professors and students are unaware.

The structural fact that most AI development is being done by a quasi-monopoly of a handful of tech giants raises the question of who has access to student and teacher data, for what purposes that data is used, and whose interests are being served. Universities face the concerning possibility that a small number of commercially-oriented corporations may actually set educational agendas in non-commercial institutions of higher education through the technical choices embedded in the algorithmic design of their products.

A second, well-known concern about Big Data reliant AI systems is these systems can embody, and further entrench, **pernicious biases**. For example, some widely used risk AI software – e.g., for predicting an individual's likelihood of committing future crimes (e.g., the COMPAS algorithm), or creditworthiness, or future work performance - have been found to exhibit racial biases, even if the systems do not have explicit racial categories in their system of classification and training. Often these algorithmic biases are difficult to spot, due to the massive complexity and opacity of deep learning models, or to legal protections that allow companies to keep their algorithms a "trade secret." Similar concerns about latent algorithmic bias carry over into education.

Because Big Data AI systems use vast amounts of past data to profile students, determine their preferences and make predictions about them, there is a significant risk of the intake and transmission of biased information. Built-in biases might lead to the perpetuation of past inequalities and to unfair predictions of future failure for students who performed poorly in the past. The embedded dataset bias – stemming from the original information fed to the

algorithm – can rapidly turn into a series of other categories of bias which are generated by the automatic decisions taken by the AI, in a vicious circle.

Generative AI systems have shown promise as **personalized tutoring systems**, to complement university teachers—as with the Khan Academy’s Khanmigo virtual tutor, based on the GPT-4 system. Such systems can spot when students make mistakes in math questions and help them see what they got wrong and guide them to the correct answer, in an individualized fashion. Chat-GPT can help teachers prepare lesson plans and course syllabi. Khanmigo and other generative AI systems can allow students to “converse” (virtually) with chatbot simulacra of historical figures, based on troves of text.

AI programs may be able to **cut educational costs** by automating tasks performed by large numbers of humans. For the same reason, these systems give rise to another worry: the future of (un)employment of professors. Breakthroughs like GPT-4 raise “**extinction risk fears**” in a wide range of fields, and higher education may be among them.

The possible widespread use of AI tools in the higher education sector carries also a **risk of dehumanization of the educational process**. As the role of the professor undergoes significant alterations with the use of AI, the experience of the students will also be considerably altered. Potential worries include: alienation from the educational goals of fostering deep insight and critical thinking (brought about by an increasing dependency on digital, always accessible information) and emotional detachment from the lack of embodied presence and socialization with peers and professors.

Universities in general—and Catholic universities in particular—aim to form students into thoughtful and discerning human beings capa-

ble of engaging in sustained reflection on important human issues. Unfortunately, as any college professor knows, students (and sometimes professors!) must fight a difficult battle against new forms of **distraction and inattention fostered by addictive digital technologies**, often powered by AI. Algorithms used by many social media platforms are trained to capture users’ attention and maximize time on the platform and “click through” (the likelihood that users will see and click on paid content). Numerous studies have shown the deleterious effect on our powers of attention and thought by intensive use of social media platforms with endless “scrolls” of hyper-palatable digital content. Many users themselves report concerns about their fragmented attention span and difficulty managing their relationship with addictive digital technologies. In this techno-social environment, universities must reflect on how they can foster practices of attentional focus, concerted thinking, and virtues of self-control in resistance to tech addiction.

At a more general level, these recent changes will also have important consequences in an already unequal academic world. While some have argued that the recent technological developments leading up to AI tools have promised to bridge significant inequalities in education – for instance, by scaling up personal tutoring to a larger number of students, or making available a number of tools often reserved to well-funded institutions – these systems are far from unproblematic from the vantage of **accessibility and equality**. Beyond possibly contributing to “mask” underlying inequalities in access to education – by serving as possible excuses for policymakers not to tackle underlying, structural issues – the disparities in access to these state-of-the-art systems can further widen the social and economic gaps already existing at a global scale.

- If we understand education to be a process of intellectual and moral *formation*, even transformation, rather than the

mere transfer of information or development of skills, how will that shape our sense of the promise and perils of AI?

- What will be the role of professors and educators in an age of widespread use of generative AI tools that are seemingly capable of transmitting and summarizing information, turning it into “pre-packed parcels of knowledge?”
- How will the learning process – hitherto centered on the classroom and face-to-face discussion with peers and professors – evolve, as AI systems, which are mostly impersonal and do not require direct, human discussion, are increasingly present in students’ formation?
- How can universities strengthen students’ (and teachers’) powers of attention, deep focus, and concerted thought, when we live surrounded by AI-powered digital technologies that are both essential to modern life and highly addictive and distracting (part of the “attention economy”)?
- In what ways might AI foster (or undermine) our universities’ aspirations to be communities where people can encounter the transcendent and spiritual dimensions of reality? How might these technologies, and the attitudes that go along with them, shape our engagement with God and with the sacred?

III Research and Scholarship

Obviously, AI is already a **topic** of much research and scholarship. This includes work being done by computer scientists, mathematicians, psychologists, linguists and others to create and develop AI systems. It also includes

work by philosophers, sociologists, economists, historians and others to understand the full meaning and significance of AI. It is safe to say that the amount of scholarship devoted to AI is only going to increase in the near future.

What about the **use** of AI in research and scholarship? Emerging AI systems are introducing a new dimension to the long-standing debate on **academic integrity, plagiarism, and academic fraud**, brought on by the newfound possibility to generate sophisticated texts able to elude most plagiarism detection tools. While most experts agree that these tools can be integrated into both the educational and research processes of the university, there are several ethical challenges surrounding their application. In the near future, beyond plagiarism, we might also face the dissemination of unchecked inaccuracies and “digital hallucinations” of generative AI – which can undermine scientific soundness – leading to the self-perpetuating “automation of error.” Generative AI systems can “hallucinate” articles that do not exist, fabricate historical “facts,” and even make up potentially damaging falsehoods about individual’s background.

AI tools can also effectively serve as vehicles to intentional **misinformation** at an unprecedented scale and speed. The form and extent of this potential (mis)use of AI systems is hard to predict, and prominent academics and industry researchers (including Danielle Allan, Geoffrey Hinton, Kate Crawford, and others) have warned of the potential of AI to flood the information sphere of democratic societies with false, harmful, and toxic content—even at a mass-scale, personalized level—and thereby undermine public discourse and damage democracy. Armies of chatbots could target citizens directly. The ethical risks here extend well beyond the educational sphere, and map directly onto pressing concerns regarding the relation between information, political participation, and democratic freedom. But the university clearly has an important role to both spread awareness of these dangers—through research

on the issue—as well as to form students in such a way that they themselves can exercise “digital discrimination,” thinking critically and carefully in a muddled and manipulated information environment.

The current world of academic research is centered on the publication of research outputs – the “publish-or-perish” logic. This intersects with AI in concerning ways. The mass-production centered logic of research will likely be reinforced by the introduction of generative AI tools, as the latter are integrated in the research process. As chatbots are gradually entering the actual research process – with students and researchers making use of its generative capabilities to write parts of articles and essays, especially those involving literature reviews and contextualization – some scholars have raised the possibility of a significant **increase in publication numbers**, which would not be accompanied by a real increase in expertise, knowledge, or quality. To address the “publish or perish” culture and the replicability crisis, a prudent strategy involves encouraging or mandating researchers to disclose their work transparently when submitting scientific contributions. Some journals have started incorporating specific questions during the submission process (e.g., Springer, Elsevier), requiring authors to explicitly disclose if any part of their manuscript was generated with the assistance of AI. This emerging practice reflects the growing awareness of AI’s role in research and aims to maintain transparency in the academic publication landscape. These worries have already gained special prominence given the replication crises in the social sciences. If generative AI leads to even higher demands for quantity and speed of research output, even setting aside worries about “digital hallucinations,” will we see a further decline in academic quality? Beyond the immediate effect of a **distortion in the professional academic environment**, this would strengthen an already heavily skewed intellectual environment in which the deeper, “slow thinking” of the humanities loses

ground against the highly-specialized, narrow, rapid-production academic output.

Catholic universities can also serve as a bulwark against the publish-or-perish incentives that drive the potential misuse of AI tools in research. Catholic universities are traditionally a place of deep, contemplative thought, where the focus is on thinking the best thoughts, rather than rushing to publish the most papers. By emphasising and rebuilding this tradition, Catholic universities can provide a necessary contrast to the rush to publish mentality, and a place in which the incentives to misuse AI tools are largely absent, thereby providing a foundation for a richer and more responsible research culture.

- How will AI systems change our current model of professional academic research? Will AI reinforce the existing high-quantity (and sometimes low-quality) publication-intensive model? Or, will we see a turn towards more critically inspired forms of inquiry, in an effort to emphasize the gap between the content that can be automatically generated by AI tools and the “deeper” knowledge rooted in human understanding, and social-relational and critical processes of thinking?
- How will different university disciplines respond to the potential automation of more parcels of the research process? Are the humanities, in particular, better suited—or worse suited—to make good use of AI tools, relative to the natural sciences?
- Can the use of AI systems serve the purpose of furthering the public availability of knowledge produced in academic settings, or will these tools widen the gap between highly-specialized, academic-oriented information,

and thoughtful, philosophically relevant knowledge.

- Which aspects of AI research should Catholic research university prioritize? What, if anything, in the Catholic tradition is especially relevant to researchers and scholars who are devoting their attention to AI?

IV The University and Society

Most contemporary universities publicly maintain that their *raison d'être* goes beyond the mere training and credentialing of future professionals for the labor market. However, with regard to undergraduate education, many universities are, in practice, oriented towards this instrumental role, and many justifications of the “value” of specific disciplines and degrees rest on improved job prospects. In this, universities are not that different from other credentialing agencies and vocational training programs. As institutions of higher education exist within the constraints of the current social world, it would be unreasonable to think that they could prosper without providing their students with tangible opportunities to promising careers. However, if we are to seriously consider their role in a more holistic, community-oriented fashion, it is crucial to reflect on their general *social* purpose, beyond supplying skilled manpower for economic growth. This is especially true of Catholic universities, given their stated missions. The recent introduction of AI tools in the academic world provides an important opportunity to articulate and critically assess the distinctive, non-instrumental, non-labor-credentialing function of universities—that is, what **intrinsic human goods and values** are sought by universities in general, and Catholic universities in particular. How might AI and related technologies play a part in this aspect of the university?

If, especially in the last century or so, academic institutions have become increasingly segmented and gradually more specialized in particular domains of enquiry, one could start by asking if this route – one that mimics the trajectory of other strictly business-oriented organizations – has not led the university to stray away from a more encompassing role in providing a more unifying and holistic understanding of the social and natural worlds.

Perhaps the disruptive challenge of AI can encourage us to re-conceive of the university as a distinctive place, set apart from the instrumental logic of the world of commerce, whose task in society goes beyond the instrumental purpose of training students to perform particular economic activities. In that regard, might AI tools allow for a reorientation of the role of universities - especially Catholic universities - in modern, high-tech societies? This requires a more subtle response to AI either simplistic techno-optimism or pessimism. If we aim to articulate and reorient the values, virtues, and skills to be cultivated in an institution of Catholic higher education, we could start by assessing the ways in which advanced technologies like AI can further, or hinder, human cultivation.

Many ethicists, analysts of technology, and government regulators, are insisting on the urgent need for “human-centered” AI, and AI that “aligns” with “human values.” But what does “human-centered” technology really mean? And what are “human values”? Such questions are difficult for late modern societies with cultures that are largely secular and materialist. Catholic universities can play a crucial role in this debate by drawing on a long tradition of thought about human beings as social and spiritual creatures with a unique value and dignity, a responsibility for others, and an orientation to what is true, good, and beautiful.

Finally, modern intellectual culture has a longstanding strain of anti-humanism and mechanism: the view that human beings are simply

fleshy machines, hence lacking any special significance or value. The view that human beings are simply machines, their minds being “software” running on the deterministic “hardware” of the brains, has been closely associated with AI since its inception. In a memorable phrase the AI pioneer Marvin Minsky called humans “meat machines,” following in the footsteps of La Mettrie and Hobbes. In the context of AI advances, this antihumanism has ethical significance: humans can be seen as mere steppingstones to a superior form of non-biological machine intelligence.

The prominent transhumanist AI and robotics researcher at Carnegie Mellon, Hans Moravec, has argued that AI-powered robots are our inevitable evolutionary heirs and replacements—our “Mind Children.” He writes: “Like the biological children of previous generations, [robots and AI] will embody humanity’s best change for a long-term future. It behooves us to give them every advantage and to bow out when we can no longer contribute.” This is a striking—though increasingly common—anti-humanist vision of the future role of technology. As technology advances, the uniquely human form of life will, like the Neanderthals, be surpassed and eventually discarded.

We see it as essential that Catholic universities provide alternative humanistic visions of the future of human beings and society alongside intelligent machines. Not human obsolescence and replacement, but human flourishing, learning, wonder, creativity, and community.

- Can introduction of new AI tools reorient our priorities in higher education, or will it only strengthen the existing logics within the contemporary university?
- Is it realistic, in the age of AI, to think of a generalized change when it comes to the existing academic model of undergraduate education, largely centered

on providing credentials and signals for the job market?

- Can the university—and Catholic universities, in particular—provide a space of deep and leisured reflection on fundamental human values in goods, thus serving as a necessary counter-weight to a world of rapid, often heedless technological change?
- How can Catholic universities articulate a humanistic alternative to anti-humanist and transhumanist visions of the future of human society alongside AI?

Conclusions

If it holds true that the Humanities play a fundamental role in analyzing and addressing various aspects of AI, including ethical, legal, and philosophical issues, it is clear that AI is significantly impacting the Humanities themselves. Today, the Humanities can and should evolve into an experimental field. Across the centuries, humanists have always used data in their daily activities, but nowadays AI offers them an unprecedented amount of data that can be processed automatically with a level of quality unattainable by humans. Such a data-driven computational breakthrough opens new perspectives and challenges for humanists in the era of AI, which is the era of multi/interdisciplinarity. Indeed, today a sharp separation between the Humanities and the Natural or Hard Sciences simply does not make sense anymore. Multi/interdisciplinarity must also impact education, with Humanities’ students learning mathematics, and computer science students discovering how their skills can be applied to the Humanities. We are transitioning from the concept of Digital Humanities, which has been

winning over the last two decades (Schreibman, Siemens & Unsworth, 2004)⁷ to that of a unified Humanities in a digital and computational world ruled by AI.

However, emphasizing the role of AI and computation in the Humanities does not mean to diminish that of human beings. We must always remember that AI has no inherent consciousness or ethics. Nevertheless, we tend to attribute human qualities to it, talking about what it “says” or “thinks”. This poses the risk of misunderstanding its fundamental nature, which is that of a set of mathematical techniques that perform tasks associated with humans without awareness of meaning. The challenge lies in shaping human ethics to guide the development of AI ethics, as well as in exploiting at best the contribution of AI in the Humanities to study and/or perform human creativity. As a matter of fact, AI can generate seemingly creative results, but it is, in reality, just based on statistical generalizations built upon (very large) training data. However, AI can identify data correlations that escape human observation, opening up new opportunities for data processing: exploiting such correlations hidden to the human eye represents an added value and a challenge for the Humanities at the dawn of the era of AI. Also, the hidden nature of AI systems opens another crucial issue, namely the explainability and transparency of AI. Many AI-related resources are like “black boxes,” capable of making predictions but challenging to decipher. It is essential to develop more transparent AI systems to understand how they make decisions and which data influence these decisions.

The role of universities, especially Catholic universities, is crucial in addressing all these challenges. They must inform and guide the evolution of AI toward a more human-centered approach, respecting human dignity and avoiding delegating moral responsibility to machines.

This requires an multi/interdisciplinary approach and active engagement in education about AI and ethics in the development and consumption of AI. We are living in an exciting but demanding time, primarily due to the very nature and role of the Humanities and, more broadly, of humanity itself.

⁷ Schreibman, S., Siemens, R., & Unsworth, J. (Eds.). (2008). *A companion to digital humanities*. Wiley & Blackwell.