

The Ethics of LLMs at Universities: A Case for Restriction and Regulation

István Zoltán Zárdai

Educational Lecturer at Juntendo University

Visiting Researcher at Keio University

‘Disruptive technologies’ is a euphemism for new technologies released lacking adequate regulation, causing significant unemployment and costly, inefficient additional labour. So it stands with LLMs. They output lookalikes of authored writing. Most output remixes existing materials, effectively stealing, since lacking understanding and intention original meaning is not added.

LLMs enable low-cost, high-reward dishonesty. Students attempt to submit these products as their own texts. Some in education propose to use LLMs to allow students to generate text and then revise it. This is feasible in groups small enough to ensure that students can be monitored continuously and substantially rewrite the texts while reflecting on the editing process. However, such highlighting of what is important about rewriting tasks reveals that LLMs do not add value to such activities.

The recommended policy is to restrict and regulate LLMs. This protects jobs; it ensures that students have to develop skills that enhance their agency and their character development; and alleviates the problem of massive scale intellectual property misuse.

Resistance to AI is urgent: no company promoting it will back down from doing so, no matter the risks of their products. This is a case of a game where one agent will never cooperate, and both agents endorse incompatible value sets. Consequently, no group potentially adversely affected should have qualms about intense opposition and has a duty to lobby for regulation and conditional supervised rollout of LLMs.

1. Introduction

The main points of this paper are that, first, Large Language Models (LLMs) are massively harmful in numerous ways. Second, that we need to ensure that appropriate regulations and controls are in place, especially in academic contexts where LLMs enable low-risk/high-reward dishonesty. And finally, that many of the motivations for adopting LLMs seem to rest on mistaken criticism of higher education (HE) and particularly teaching quality at universities. Instead of endorsing such tendentious and factually incorrect stereotypes we should focus on asking what we teachers can do in the age of AI and how can we do it best. We can come to realise that higher education’s role is not changed by LLMs. Regulation and prudent limited use for beneficial educational purposes is what we should engage in.

The paper argues for these claims by explaining in section 2 relevant important aspects of LLMs. Section 3 discusses HE specific problems with LLMs, especially negative stereotypes pertaining to the educational quality of HE institutions. Section 4 turns to questions that the academic community needs to work on more, and finally in section 5 policies and regulations are recommended that we, as the responsible community of HE stakeholders, should push for.

What the paper does *not* argue for is that we should ban or stop using LLMs entirely. If we can use new technology safely we should, especially if it makes administration, planning, grading, and teaching easier, more efficient, and helpful for ourselves, our students and society. If LLMs are used well, in restricted ways, they can enhance learning while also providing educators and learners another useful skillset.

2. What is important about LLMs for the purposes of this paper?

LLMs are machines detecting statistical correlations - programmed by humans, in line with company objectives - to generate text. They are self-supervised learning models that predict tokens (language units). Their performance is fine-tuned by supervised learning and reinforcement learning. Reinforcement learning needs extensive human feedback, it is not something any AI or other robot can do, and is effective because of the large number of humans who put in enormous effort and working hours. Companies building LLMs install filters to limit harmful outputs and to prompt LLMs to aim at providing conforming replies, so as not to upset customers and avoid negative news coverage, shielding their investors and their company.

Some impressive relatively recent achievements that enabled current LLMs, (like ChatGPT, Perplexity, and Bard) to reach new levels of performance are that LLMs are now training each other, and that they can call on other apps and plugins to supplement their weaknesses. The classic example of this is ChatGPT, which cannot count, hence uses the calculator app inbuilt in Windows.

3. Negative stereotypes concerning higher education motivating LLM use

There are several reasons why certain companies, politicians, managers, and researchers urge rolling out LLMs rapidly and in masse. The quality and truth-aptness of these reasons varies widely. Let us briefly review some of the more prominent ones. There is an atmosphere of urgency, built up by mass media, suggesting that HE is late to something important. Why is there such a general atmosphere?

An obvious reason is the promise of benefits to institutions, workers, and students, as well as society at large. People are excited about opportunities, and that is a normal thing. There has been a shift in HE during the past decades with more and more governments – foremostly in Australia, the US, and the UK – pushing HE to mimic competitive, profit-driven business sectors. This led to a change of perspective in HE with managers adopting a competitive mindset and trying to gain advantages by adopting technologies early, taking on risks to potentially gain an edge, with little regard to the existence of evidence or data supporting their efficiency and overall beneficial nature. This climate of urgency and haste makes many people forget that in reality it depends entirely on HE teachers how and how much they use LLMs, and how much they allow their students to do so. Doing more faster rarely benefits quality, and since the rollout of LLMs in education depends entirely on HE teachers, the profession should first consider possible regulations and changes to LLMs based on their risks and their harmful impact, before deploying them. A slower, more restricted rollout – for example in the form of LLMs designed and appropriately restricted for the purposes of HE – might be much more beneficial for, among other factors, learning outcomes, academic integrity, student upskilling, classroom management, mastering IT skills.

HE leaders, administrators and teachers also have to keep in mind that we have a Duty to Protect: we should not threaten the wellbeing and safety of our students, and we should aim at promoting safety, equality and quality in HE, as well as in society at large. Education is a public good that serves communities primarily, and individuals second. As such we have to keep in mind the impact our adoption of new technologies has on society, both in terms of what HE means to others, and in terms of actual impact on the environment, businesses, society, and unemployment. Again, what is argued here is not that LLMs should be banned. However, more studies, experiments, and trials on the effects of LLMs on students, on academic workers, and the wider society are needed.

Some of the pressure to rapidly adopt and deploy LLMs is fuelled by popular misconceptions about higher education. There are several false, unverified and contestable claims parroted by private corporations, politicians and mass media about education in general, and more specifically about HE. Five of these claims are presented and evaluated here briefly to further contextualise and defuse the unfounded techno-optimism present in the current mainstream discourse.

The first popular trope is that HE is in a crisis, and in large part this is so because teaching quality is not good enough. The truth is that the quality of education – while it should continuously be improved and developed – is remarkably high compared to any previous stage in history. Teachers are well-trained, productive and their labour is very high value, socially constructive, especially in comparison with fields like finance and marketing, two popular fields attracting large numbers of graduates. A related and similarly unbased concern is that education is failing students, graduates perform poorly at work. The reality is that worker productivity has been improving significantly and gradually since World War II. Today's workers, including recent graduates perform much better and get more done in the same amount of time than workers did 80, 60, 40, or even just 20 years ago (Economic Policy Institute 2022). While it is true that this does not translate into fast economic growth, this seems to be due more to capital underperforming, rather than worker efficiency not being suitable. At the same time real wages are down: workers today get more done, but are paid less, their purchasing power is lower than their parents', as is well documented by several large scale economic studies from the past 10 years. If schooling has an effect on workers efficiency, then HE seems to be contributing to this constant increase, and there is no evidence to assume the contrary.

The second popular mistaken claim is that student satisfaction and gains are true reflections of, and depend mostly on teacher performance. Some governments – for example in Britain – have been implementing quality assessment systems in which student evaluations feature largely and poor evaluation carries great weight. These attempts at improving student outcomes and experiences are to some extent well-meant, and have beneficial outcomes; nevertheless, they lack significant evidential support. Large scale studies found that student evaluations are informative but suffer from several biases. Just to mention some significant examples: students rate male teachers higher in general; students rate teachers perceived as funny higher in general; students rate older teacher higher in general. For detailed recent discussions of issues with student evaluations see Esarey and Valdes 2000, Kreitzer and Sweet-Cushman 2021, and Aloyo 2023. It is also established that student gains depend primarily on what levels of knowledge and ability students entered university with. This in turn depends to a large extent on the family background of students and is tied most significantly to the educational background and income levels of their parents (Dubow, Boxer and Huesmann 2009).

A third prominent idea is that HE is not sustainable. However, even if this point would be true, the main reason for this is not due to HE's performance. Public funding for HE has been massively curtailed during the last 30 years due to political decisions, many of which might be attributed to corporate pressure and to harmful, extremist market-ideologies. A connected claim is that to overcome such sustainability issues performance improvements should be implemented as a priority.

The former trope is coupled with the unsubstantiated claim that new IT products overwhelmingly lead to productivity and quality increases, and as such disruption is not only something we should tolerate, it is beneficial. This is not born out either by data or studies specific to HE, nor by analogy from other sectors. Recent research in economics seems to indicate rather that the effect of disruptive new technologies depends on the specifics of given national sectors, and disruptive innovations and organisations using them can have negative effects on overall productivity (Feder 2018). What we see from increased deployment of recent platform economy tools and other forms of automation is increased unemployment, worse services – not always accompanied by price decreases –, unwillingness on the part of IT companies to adopt the policy suggestions of regulators, rise in exploitative work, and many tasks that were offered as parts of services – booking a taxi through a dispatcher – being offloaded to users of services – people have to download and learn to use new apps for almost every service, many of these apps collect private data invasively, and require users to take several actions that normally a company employee would perform.

In truth, so far new IT products have not contributed significantly to productivity in HE, and even in fields where it did – for example in finance – it took decades to do so. Benefits thus seem to be incremental and take long time to appear. Educational technology in specific has not impressed: there have been thousands of IT products for HE since the early 1990s and very few have proven to be actually impactful, useful applications that are widely adopted (Syverson 2017).

It can be concluded then that most criticism of HE teaching quality seems to lack substantive support. The reality is that HE seems to perform in-line with the overall economy, constantly developing and training productive, ever more highly qualified workers. Even if HE would not have any other social uses and purposes – most surely it does, just to mention as examples research, researcher training, knowledge dissemination and character development of citizens – it seems that it performs exceptionally well. The takeaway is that despite oft-repeated phrases there is no need to rush with the deployment and unregulated use of LLMs or other technologies in HE, or education in general. In fact, regulation has proven to be effective and producing just outcomes almost every time a serious effort at it had been made; for some examples on successful regulation see Oreskes and Conway 2010: 105-6, 162-68, 236-39, and for some constructive proposals on regulating the global gig economy that employs AI and digital platforms see Wood, Graham and Anwar 2020. Wood, Graham and Anwar also note that the global gig economy is actually operated largely by five countries and a handful of companies. The same seems to be the current situation regarding LLMs, as well as AI in general. Hence, regulation is in fact much easier than is usually presented in public discourse.

4 Three types of HE specific issues with LLMs

After reviewing the general issues with LLMs that might motivate us to slow down and regulate their rollout, I mention three groups of HE specific problems with LLMs. The first of

these concerns job security and stability in HE, the second is the many serious issues with academic integrity, the third is the deskilling of students. While everyone in HE is familiar with a simple statement of the problems, the paper argues that we need to interpret them more specifically and some solutions to them already exist that we can make use of. Due to the paper's focus I sidestep concerns with employment and stability of job-roles, and survey the two groups of problems that are most closely education, academic honesty, and skill acquisition related.

4.1 Challenges to academic integrity

Whenever the topics of academic integrity and honesty come up, we have to reflect on this: what is the purpose of teaching in HE? For example, should a computer science teacher teach coding, or how to prompt the LLM to build a code, and then how to prompt it again and again to make it more elegant, effective, and fool-proof? Should an academic English teacher teach students how to communicate in English? Or how to prompt an LLM to do a specific task well, for example, how to translate the right parts of a text, or how to paraphrase something in the right tone, voice and register? Should a philosophy teacher teach philosophy – content and methods, like critical thinking, aiming at rational equilibrium, devising thought experiments, and so on – or how to use an LLM to come up with replies to questions, and how to make and defuse arguments?

These questions are genuine: if a teacher teaches their students to use LLMs to do their tasks, the students do not acquire the relevant language or reasoning skills and abilities. This prompts reflection on what teachers in HE should aim to achieve with students. Regarding this the following are reasonable assumptions: teachers in HE aim to increase the employability of their students; they aim at helping them to become a decent, well-informed citizen of their society; and to acquire the abilities and skills that receiving their diploma is usually assumed to be conditional on. This includes certain work ethics, field specific skills, and some general cognitive abilities.

The interesting question then becomes, whether using LLMs threatens the acquisition of these abilities and the achievement of these goals? The realistic answer is that this depends on how it is used and regulated. Allowing students to take shortcuts by using LLMs to circumvent tasks that would provide them a useful stretch harms their development. Also, eliminating human face-to-face interaction between teachers and students threatens to affect their character and social skill development adversely. Hence, it is reasonable to recommend avoiding reducing teacher numbers – if anything increases to achieve better teacher-to-student ratios would be beneficial – and using software based on LLMs that is specifically adjusted for HE needs. This could mean software which only answers a limited set of prompts – for example, it can provide suggestions, but it is integrated with submission software to prevent or to highlight the submission of the AI generated content; or giving teachers the option of sharing the material with the software which then does not allow students to prompt it for answers on tasks in the material, while providing guidance on the rules and tips on how to tackle the material on their own. Governments can incentivize companies to develop such products by providing funding sources or grants to HE and other educational organisation for obtaining them; as well as by regulating how LLMs can be made available and accessed online, making it difficult to access them without that being obvious to teachers, proctors, and other supervisors.

These proposals lead us back to academic integrity. LLMs, as they are implemented and made accessible currently support, enable, and commit dishonesty. It would be reasonable to make the companies making them available as they are now liable for some of the harms caused, including additional working hours – a precious resource for any organisation – and development lost out on. However, one might ask the question: are LLMs any worse or more harmful than alternative existing services that enable and facilitate cheating, like essay-mills?

The straightforward answer is that yes, LLMs are worse for academic integrity. They are much easier to use, they are free, their use is hard to track and prevent, they make supervision harder and detection more difficult. As the UK's QAA notes in their recent report on LLMs (QAA 2023) there is no reliable detection software, and there are serious data protection and consent concerns, especially when assessing student work. (Perkins, Roe, Postma, McGaughran and Hickerson 2023). To give another illustration of problems, work in a classroom where using computers is allowed requires supervision, which in turn presupposes that classes are small enough. It works best if students are intelligent and motivated, and have integrity. The reality is that these conditions do not always obtain. The next question is then whether this is a problem with LLMs? or a more general issue? Certainly, it is a more general issue; however, LLMs do pose it in new practical ways that call for novel solutions. And they pose a paradox: it is easier to catch out LLM-misuse if the students work on computers, but working on a computer is what enables misusing LLMs in the first place.

Preventing and detecting LLM use is thus needed – at the current stage of their implementation – to avoid misuse. This can be done in various ways. If the focus is on content, and student numbers make it feasible, oral examinations can be suitable for many subjects. Paper based tests – even just quick level-checking tests at the end of a class – can work to gauge the students' real levels of knowledge and skill, and find out whether their submissions have been artificially produced or enhanced. Also, institutions can purchase software and install it on computers which stop students from switching between applications, giving IT and teachers control over turning the software on and off. There are some such applications already on the market and some computer-based language exams use such systems as well.

Preventing LLM misuse is harder in writing intensive and language focused courses. Still, there are ways to detect the usage of LLMs. Most current LLMs use a set of typical phrases (an example: ChatGPT's message declaring that they only know about a topic up to the date of their last implementation), recognisable style (neutral, mainstream politically, noncommitted on debatable or open issues); as well as data stamps and other traces of having copied in large amounts of text at once into a document (for example teachers can check the editing history of Google Documents).

The fact that such preventive and defensive measures are needed indicates the limited usefulness of LLMs. Not to mention that with time LLMs will evolve to be harder to identify based on the recurring text patterns they employ, and students – and others expected to produce original work – will become savvier at avoiding detection. Some in education propose to use LLMs to allow students to generate a text with an LLM and then revise it (Eke 2023). If students do this on their own, it is questionable how many will actually revise the text, rather than just submitting it as it is. Using LLMs is feasible in groups small enough to ensure that students can be monitored continuously and substantially rewrite the texts while reflecting on the editing process. However, such highlighting of what is important about rewriting tasks reveals that LLMs do not add value to such activities. We could just use any paragraph from a textbook, article, website to do the same task. In fact, using texts that teachers know well and

having all students use the same text can be beneficial: it is clearer what counts as improvement, what students can learn from the task, and it is much more efficient to evaluate the work produced by the students.

LLMs pose several further issues when it comes to assessments. A typical academic writing task would include most of the following steps: choosing a topic, formulating a claim, finding resources, writing a draft, revising and expanding the draft based on feedback. LLMs can be used for any of these steps. If unchecked, a student could fake it all the way through. Due to this, if educators still employ academic tasks involving any of the steps mentioned, in most cases the tasks will need to be supplemented by means of ensuring that the students actually completed the steps themselves. This can be done via oral presentations, short in-class writing tasks or tests, additional tasks that indicate the students' actual ability to write complex prose. A method that the author has tried and worked reasonably well was to supplement essay submissions in a Philosophy seminar with presentations, debate sessions, tasks where students had to evaluate arguments (valid or invalid, sound or unsound, true or false, supplement missing premises, make hidden assumptions explicit, identify ambiguity, and so on), and they had to attend a one-on-one consultation to discuss their first drafts. Such methods are of course practical only when student numbers in a course, and in general teacher-to-student ratios permit it.

These considerations also make it clear that regulations have to fit the goals of the programme to which they are applied. In *content-based* classes – like geography, philosophy, accounting – it is fine to use an LLM to help improve the grammar or to paraphrase a descriptive paragraph. However, it is problematic to let the LLMs suggest results, fill in steps in an argument, or to work out the conclusions instead of students.

In a *language-focused course* it would be fine to use the LLM to pick a topic and generate text suitable for assignments. Still, it would often defeat the purpose if students would be allowed to let the LLM do the work on paraphrasing, finding grammatical problems and fixing them, fine-tuning the tone, and so on. This does not mean that advanced students who already possess strong English, say, upper-B2 / C1 level, and have good study and self-regulation skills cannot benefit from supplementing their study by using LLMs to revise their work, suggest improvements, and so on. But such students are the minority of students who take courses in English, English language, or academic writing, since undergraduates often do not have the right-level of independent study skills yet.

All this suggests that it is fitting to have more general and broader-brush regulations at the institutional level. The benefit of this is that it can leave enough room to interpret and apply the regulations as institutional investigations find it most fitting (within the bounds of fairness of course). At the programme level more detailed regulations are needed most likely, and programme members are in the best position to decide what regulations would create the right mix of encouraging and enabling beneficial use, while preventing and discouraging harmful use. Currently the trend is that institutions try to shift the burden of coming up with reasonable policies to individual teachers. There is no reason why departments, centres, or programmes could not make their regulations, and with the support of institutions make use of the disciplinary framework of the organisation to render regulations effective.

4.2 Worries about deskilling of students

The third group of higher education specific issues consists of potential risks of deskilling: of not developing or of developing important skills at sub-par levels. This worry has three sources. The first of these is *cognitive laziness*. This means the temptation to outsource tasks involving reflection, deliberation, use of empathy, imagination, learning to make conceptual distinctions, weighing reasons, and related tasks to LLMs. This is problematic both for the reason that learners fail to develop these skills adequately, and also because LLMs cannot engage in any of these processes, they simply mimic or steal text that was produced by humans who engaged in such processes, leading to texts of dubious quality and moral status.

At this point intellectual honesty requires us to investigate whether this worry is well founded or not. We can ask whether the same has not occurred earlier with the introduction of such historic technologies as printing, tv, and the internet. The reply is that the changes brought about by these technologies are similar, but not identical. One of the main reasons is that previous technologies required substantially more human input, both in creating products, and in coming up with the content. Another important difference is that the mentioned technologies were mainly aimed at disseminating existing content, whereas LLMs create new content when prompted. Furthermore, due to the ubiquity of computers and phones for completing tasks in class, the level of presence and ease of use of LLMs is new. Cheating by using a book or notes is possible but certainly harder given appropriate proctoring. An example is the widespread use of Google Classrooms. Educators endorse the technology partly because it makes grading and providing feedback much faster and possible from anywhere. Alas, if 25 students are working on a task online simultaneously, it will be very hard without additional software or proctoring, for one teacher to ensure that none of them use LLMs to generate submissions. Hence, it seems realistic to worry about deskilling.

Another question we have to ask at this point is whether deskilling really is an issue. After all, how do we know what skills will be needed in the future? Perhaps some of the skills that students do not acquire at the same level as before will be less important due to the prevalence of generative AI in the workplace as well as in the life of citizens, and our private lives. Will coding or copywriting skills be important? Perhaps, but perhaps they won't. The honest reply is that we simply do not know what exact skills will be useful. This is not a failure of HE: governments, leading companies, ministries, technology experts – while some of them are confidently outspoken – seem to be in the dark as well, offering little specific guidance, seeming to try and pass the buck. This seems to suggest that we might as well just allow LLM use in HE.

However, this would be a hasty and mistaken conclusion. LLMs should be used sparingly and be subject to strict regulation. Besides enabling students to gain a foundation in specialist fields like law, economics, mathematics, biology, or accounting, HE has historically also done a good job at equipping people with general skills that are not field specific, yet contribute both to personal wellbeing, peoples' ability to be responsible members of their community, and to make them able to learn new skills, and seek out and acquire new knowledge so that they can keep developing. HE is good at instilling these general skills through teaching a specific body of knowledge. While legal reasoning is different from reasoning in finance or in engineering, there are some general reasoning skills – for example epistemic skills about assessing the probability of the trustworthiness of different sources of information – that can be imparted through all university courses. This function of HE is valuable and important for society: it

contributes to a liquid job market and to the prevention of alienation at the same time, while increasing the quality of public and intellectual life. Based on this we can say that there is no strong reason for HE to simply follow trends and uncritically endorse LLMs, or to try and provide the training in using new technology that employers should provide. In fact, there are substantial reasons to retain its mission and continue doing well what it is successful at. Writing is an excellent task because it involves many skills – but in the end the improvement of the skills is the goal, not writing itself, and the skills can be improved while trying to prepare for something, especially if thinking and reasoning are involved.

5 Questions that need to be addressed in the future

The discussion of LLMs risks should not lead us to endorse or promote a complete ban. This is not in the interest of HE stakeholders. Educators, researchers, managers, and regulators should provide funding and time to explore beneficial teaching and research uses. HE stakeholders need more education that helps them understand how LLMs work, instead of endorsing the popular overstatements echoed by corporate-funded media pieces. It is important to understand why LLMs perform so efficiently at some tasks, but are so inefficient at other – seemingly – similar tasks.

5.1 Established harms in HE

Outsourcing more work to LLMs means that students receive even less attention from teachers, and also the danger of teachers having to teach higher numbers students because of the false presupposition that simply introducing LLMs into a workplace will make work simpler, faster, or easier. LLMs are already used widely for producing propaganda, racist, sexist, and other harmful and hurtful content. An example is how Republican legislators in the US used LLMs to add books to the list of banned titles in some jurisdictions (Williamson 2023). Using the same products without any filters, training or protective measures means we are exposing our students – and in some cases paying customers – to propaganda, misogyny, racism, military disinformation, politically motivated fake news, and more.

5.2 What exactly do LLMs contribute to classroom work?

All these concerns lead us to ask what contribution LLMs can actually make. Some people discuss them in terms of ‘ban or no ban.’ Others are more flexible and thinking about how to retain the most crucial educational components of their courses while making room for reasonable LLM use (Mitchell-Yellin 2022). People who are unreasonably sceptical of regulations or overwhelmingly optimistic about the benefits of AI propose rapid adoption, and giving up on the hope for any regulations (Gleason 2022). This is an oversimplification and a question that is not helpful. The interesting question is how LLMs can be modified and restricted so as to mould them into positive forces for society and more specifically education and HE.

More nuanced approaches try to assess what ailed traditional assessments, whether AI and specifically LLMs can help overcome such issues, and what new challenges technology introduces into assessment and evaluation. See for example the work of Swiecki, Khosravi, Chen, Martinez-Maldonado, Lodge, Milligan, Selwyn and Gašević 2022. To continue along

similar lines, consider an example of customising a task for a group of students. In this case we could prompt an LLM by providing it the following information: the prompt, the material to customise, the earlier replies of students based on which the material needs to be improved, and the next material to cover with which the improved material has to have consistency. Providing all this data enables prompting the LLM to use the next material and design tasks that can help a student correct their most common errors.

Benefits of this process include that it is customised and it *might* meet the needs of students; however this is not certain since one cannot measure the effectiveness of the materials without testing them and seeing the students' results. The drawbacks are the same as when human teachers design new materials: in normal cases the teacher knows a good deal more than their students, certainly enough to guide them effectively, and still, if the teacher does not understand the situation, personality, classroom dynamics and motivation profiles well, a good material might not achieve the intended outcomes.

5.3 The core issue and some deliberately confusing ideas surrounding it

The main claim of this paper is that most concerns about LLMs in HE are at their core a political and a professional ethics issue. Applications are rolled out without proper regulation, testing, or consulting, and do a good deal of harm. The organisations making these products market them by mislabelling harm as 'disruption', 'revolution', 'innovation', and the push of competition. This language disguises the fact that the change is good for the seller and bad for large swathes of society. Contrary to what the disruption-language implies, if people get paid the same or even less, and have to again and again change how they work by learning to use new tools or change their job tasks, that is not a benefit, that is a burden. More narrowly, if solid forms of assessment in HE become impossible to use – like essay writing – that is a problem, not an opportunity. An opportunity would be to pay someone to have time and instead of their usual duties be allowed to focus on developing or innovating tasks and exercises. If people are forced to rapidly react to technological change in addition to their normal workload that is a difficulty, a stress factor, and a threat to performance.

The educational sector should back politicians who are willing to resist the push to capture educational budgets for tech purposes, and who are willing to legislate to prevent layoffs. Paying for access to such services can be a problem and the threat that public funds will be 'hijacked' and channelled to the IT sector to an even larger sector are real (Williamson 2023).

The quality and human centred nature of education – as well as social work, law, policing, and healthcare – need to be preserved, and in fact improved. It is very important to keep in mind that the educational sector does have the well-trained and able human resources to push these three lines of lobbying activity at the same time. As a recent UNESCO publication on AI in education noted age limits for use and amount of exposure need to be specified, with some ages not using LLMs at all (Williamson 2023). The UNESCO report also voices concerns about the misuse of student data – it is not clear how such data is processed, stored, used – and the lack of privacy while using LLMs, as well as data gathered about use by private and institutional users being sold off without robust consent or control.

For the HE sector it should also be especially important to set up regulations that protect meaningful public discourse and prevent the spread of misinformation, government, military and corporate propaganda. That is, regulations should foster equal access to good quality information, and prevent large scale exposure to biased information produced by already

powerful organisations and social agents. Most likely this requires significant mandatory public control over LLMs. Such control is justified, and it might imply that we have to think in terms of collective or distributed responsibility—as these concepts are applied in law, international relations, and philosophy. For such proposals regarding responsibility for AI see Floridi 2016. Society already experienced that social media companies are unable and unwilling to regulate themselves, and that the for-profit structures they adopt make it easy to exploit them for propaganda purposes. The fact that there have been other sources of information – theatre, cinema, books, newspapers and so on – which have been and are misused for similar purposes is not a reason, not to worry about LLMs; it is a reason to worry about LLMs and other sources of bad quality and misleading information, indoctrination, and other forms of manipulation.

Conclusion

This paper argued that it is already known that LLMs have significant social harms and also pose risks for HE. It is almost certain that organisations producing LLMs will not regulate them adequately without external pressure, while simultaneously they produce and reinforce negative stereotypes of HE and teaching quality in HE to motivate the adoption of IT consumer products.

Hence it has been recommended that HE stakeholders should cooperate extensively and as a priority with other sectors of education, as well as organisations in technology, business, citizens and rights advocacy groups, and other spheres of society, to help enact and enforce better and stricter regulations regarding the roll-out, use, and development of LLMs.

Besides enabling to refocus HE on its core missions of character and skill development, and rein in academic dishonesty, this would have the added benefits of lessening environmental harm, mass unemployment, and the further increase of the North-South inequality. More specifically in an HE context it would help protect teaching and learning quality, help enforce academic integrity, avoid deskilling, protects learners to whom we have a duty of care from unregulated harmful manipulative content, guard HE budgets from being diverted to fund IT products of questionable added value, and protect the working time and methods of HE professionals from restructuring based on external, unrelated interests.

References

- Aloyo, E. 2023. ‘Why It Is Wrong to Use Student Evaluations of Professors as a Measure of Teaching Effectiveness in Personnel Assessments: An Unjust Risk of Harm Account.’ *Public Affairs Quarterly* 37 (2):79-100.
- Dubow, E. F., Boxer, P., and L. R. Huesmann. 2009. ‘Long-term Effects of Parents' Education on Children's Educational and Occupational Success: Mediation by Family Interactions, Child Aggression, and Teenage Aspirations.’ *Merrill-Palmer quarterly (Wayne State University. Press)* 55 (3): 224–249. doi: <https://doi.org/10.1353/mpq.0.0030>
- Economic Policy Institute. 2022. ‘The Productivity-Pay Gap.’ *Economic Policy Institute*. Accessed 10 May 2024 <https://www.epi.org/productivity-pay-gap/>

- Eke, D. O. 2023 'ChatGPT and the rise of generative AI: Threat to academic integrity?' in *Journal of Responsible Technology* 13 100060?' in *Journal of Responsible Technology* 13. doi: <https://doi.org/10.1016/j.jrt.2023.100060>
- Esarey, J. and N. Valdes. 2020. 'Unbiased, reliable, and valid student evaluations can still be unfair.' *Assessment & Evaluation in Higher Education* 45 (8): 1106–1120. doi: 10.1080/02602938.2020.1724875
- Feder, C. 2018. 'The effects of disruptive innovations on productivity.' *Technological Forecasting and Social Change* 127. doi: <http://dx.doi.org/10.1016/j.techfore.2017.05.009>
- Floridi, L. 2016. 'Faultless responsibility: on the nature and allocation of moral responsibility for distributed moral actions.' In *Philosophical Transactions of the Royal Society A* 374: 20160112. doi: <http://dx.doi.org/10.1098/rsta.2016.0112>
- Gleason, N. 9 December 2022. 'ChatGPT and the Rise of AI Writers: How Should Higher Education Respond?' *Times Higher Education*. Accessed 10 May 2024
<https://www.timeshighereducation.com/campus/chatgpt-and-rise-ai-writers-how-should-higher-education-respond>
- Kreitzer, R. and J. Sweet-Cushman. 2021. 'Evaluating Student Evaluations of Teaching: a Review of Measurement and Equity Bias in SETs and Recommendations for Ethical Reform.' *Journal of Academic Ethics* 20 (1): 73-84.
- Landers, C. Forthcoming. 'ChatGPT, the CUPID Model, and Low-Stakes Writing.' In *Aapt Studies in Pedagogy*. Accessed 09 February 2024.
<https://philpapers.org/archive/LANCTC-3.pdf>
- Mitchell-Yellin, B. 2022. 'Conversation-Starter: Teaching Philosophy in an Age of Large Language Models.' *Daily Nous*. Accessed 09 February 2024.
<https://dailynous.com/2022/08/31/conversation-starter-teaching-philosophy-in-an-age-of-large-language-models-guest-post/>
- Oreskes, N. and E. M. Conway. 2010. *Merchants of Doubt*. New York: Bloomsbury Press.
- Perkins, M., Roe, J., Postma, D., McGaughran, J. and D. Hickerson. 2023. 'Game of Tones: Faculty detection of GPT-4 generated content in university assessments.' Preprint. Accessed 09/02/2024 <https://arxiv.org/pdf/2305.18081>
- Swiecki, Z., Khosravi, H., Chen, G., Martinez-Maldonado, R., Lodge, J. M., Milligan, S., Selwyn, N. and D. Gašević. 2022. 'Assessment in the age of artificial intelligence.' *Computers and Education: Artificial Intelligence*. Volume 3. doi: <https://doi.org/10.1016/j.caeai.2022.100075>.

Syverson, C. 2017. 'Challenges to Mismeasurement Explanations for the US Productivity Slowdown.' *Journal of Economic Perspectives* 31 (2): 165-86.

Wadden, P., Otsuru, S., and I. Z. Zardai. Forthcoming. 'Creating and Crafting CLIL Content Using AI Tools: A New Frontier.'

Williamson, B. 2023 'AI must be kept in check in schools.' In *UNESCO Courier* October-December 2023: 6-8.

Wood, A. J., Graham, M. and M. A. Anwar. 2020. 'Minimum Wages for Online Labor Platforms? Regulating the Global Gig Economy.' In Larsson, A. and R. Teigland (Eds.) *The digital transformation of labor: Automation, the gig economy and welfare*. London: Routledge.