Open Licensing and Generative AI for (Open) Education

Jonathan A. Poritz

jonathan@poritz.net
poritz.net/jonathan

3 October 2024

IX International Conference
UNIVERSITY LIBRARY AT A NEW STAGE OF
SOCIAL COMMUNICATIONS DEVELOPMENT
Topic 2024 "ON THE PATH TO ENHANCING EQUITY,
DIVERSITY, ACCESSIBILITY, INCLUSION"

- Thank you for the chance to speak to this wonderful group.

- Please let me apologize from the very beginning that I cannot speak Ukrainian, and thank you so much to the interpreters who do such a very hard and important job. I have acted as an informal interpreter a few times in my life, so I am aware of how hard it is to do this important job!

- It is particularly embarrassing to me that I cannot speak even a word of Ukrainian since a branch of my family comes from the area near Kyiv, and had economic and political forces been different 125 years ago, I might well have grown up there, rather than near New York city in the United States.

- I am coming to you from Italy, where I live now, working over the Internet with many open education organizations in the US and Europe, but for most of my career, I lived in the United States, where I was a professor of mathematics and computer science. I worked in the computer industry a little, mostly before I became math professor, and I used a lot of *free software* or, as I like to call it *FLOSS* – which stands for Free/Libre/Open-Source Software. So as a professor, learned of OER, which are the educational equivalent of FLOSS, and became very involved. For example, I have taught the online Creative Commons Certificate sixteen times!

- But since my background is in math and computer science, I was very excited when "AI" started to get a lot of attention in the last few years, because I thought that the open ed community would want to know more about the technology before deciding to use it, and I would be useful to that community. Instead, educators – and organizers and administrators and lawyers and politicians – seem very little interested in any of the technical questions, but rather they just want to start using it the way it has been promised to be useful.

- In our short time today, I want to talk a bit about what "AI" really is, what it really can do, and some questions around its use in OER, because of the complexities of legal rights around "AI."

- Note that my slides are openly licensed, as the icons on the lower left indicate, and are available at the URL shown also at the lower left, so you don't need to write anything down but if something is interesting, you can refer to the slides for details.

“AI” hype, part 1

We are told that “Artificial Intelligence”

- is already here
- … or will be here very soon
- is inevitable
- is a supremely valuable investment (*give me the money!)*
- makes individuals that use it far more productive than those who do not
- … so will be used in all workplaces
- … and so not teaching it would so disadvantage our students as to constitute a form of academic malpractice

- The "hype" about "AI" promises a lot, such as what is on this slide:

- is already here ... or will be here very soon

- ... or will be here very soon

- is inevitable – of course, some new technologies seemed inevitable, but were in fact harmful and so not used!

- is a supremely valuable investment (*give me the money!)* – which seems like a very self-serving thing for someone in an "AI" company to say!

- makes individuals that use it far more productive than those who do not

- ... so will be used in all workplaces – although I know many people in computer companies who have tried out "AI" for their work and quickly abandoned it!

- ... and so not teaching it would so disadvantage our students as to constitute a form of academic malpractice – although maybe teaching students actual skills would help them use "AI" in those areas more effectively, if that is at all possible!

“AI” hype, part 2

... and that AI...

- will enable everyone to become a great writer and artist
- will enable people with disabilities to have full and equal access to everything on the Internet (automatic, high-quality, detailed description of images, etc.)
- will correctly diagnose disease (highly accurate reading of X-rays!)
- ...and then cure all diseases (new miracle drugs!)
- will invent new energy sources (fusion, anyone?)
- ...and then completely “solve” the climate crisis,
- will fix homelessness and help governments plan their way past all social ills
- *i.e.*, will create a technoutopia

- More "AI" hype promises include what is on this slide, such as:

- it will enable everyone to become a great writer and artist

- it will enable people with disabilities to have full and equal access to everything on the Internet (automatic, high-quality speech-to-text, detailed description of images, etc.)

- it will correctly diagnose disease (such as with very accurate reading of X-rays!)

- ...and then cure all diseases (new miracle drugs!)

- it will invent new energy sources (fusion, anyone?)

- ...and then completely "solve" the climate crisis,

- it will fix homelessness and help governments plan their way past all social ills – the Governor of California said this a few weeks ago, which just seems *weird*....

- *i.e.*, it will create a technoutopia

“AI” hype for education

In particular, for education, we are told “AI”

- will perfectly translate any educational resource between languages, *e.g.*, English⇝Ukrainian
- will increase access to education for people with disabilities (see above)
- will act as individualized tutors, *e.g.*, in communities where there are too few teachers so class sizes are very large
- will output new educational resources when an educator merely types something like “make a 300 page calculus textbook using open pedagogy, for second-year university engineering students, in Ukrainian”
- ...or at least will output problem sets or interactive elements for ebooks (H5P) or study guides, when given a book as input and the prompt to create that ancillary resource

- In particular, for education, we are told "AI"

- will perfectly translate any educational resource between languages, for example, think of any OER you find in English being instantly translated to Ukrainian for your students to use, without taking the time, energy, and money for human translation!

- will increase access to education for people with disabilities (see above)

- will act as individualized tutors, *e.g.,* in communities where there are too few teachers so class sizes are very large

- will output new educational resources when an educator merely types something like "make a 300 page calculus textbook using open pedagogy, for second-year university engineering students, in Ukrainian"

- ...or at least will output problem sets or interactive elements for ebooks (H5P) or study guides, when given a book as input and the prompt to create that ancillary resource

The cost of "AI," part 1

All for the low, low price of

$7 Trillion

and unimpeded access to *everything ever put online*, including

- all of YouTube, Reddit, GitHub, LinkedIn, Facebook, and other social media
- every student assignment ever submitted through Turnitin and Canvas or any other LMS
- every Google Doc (even ones never shared), Microsoft Cloud document, *etc.*
- the complete text of every message ever sent or received by Gmail
- regardless of the copyright status of any of these digital objects

- All for the low, low price of $7 Trillion and full access to everything ever put online, including
- all of YouTube, Reddit, GitHub, LinkedIn, Facebook, and other social media
- every student assignment ever submitted through Turnitin and Canvas or any other LMS
- every Google Doc (even ones never shared), Microsoft Cloud document, *etc.*
- the complete text of every message ever sent or received by Gmail
- regardless of the copyright status of any of these digital objects

The real cost of "AI"

Independent estimates suggest that "AI"

- requires about half a liter of clean water per query (to cool the server farms)
- requires as much energy per query as it would take to charge a cell phone
- ... hence requiring mothballed coal-fired power plants to be returned to service
- ... and Microsoft recently announced it will pay to bring the Three-Mile Island nuclear power plant – the site of the largest nuclear accident in North America – back into service
- is losing Microsoft $20 per user per month on it's "Copilot" product
- ... and similarly for all major AI companies, so they are all losing tens of billions of dollars (at least) per year

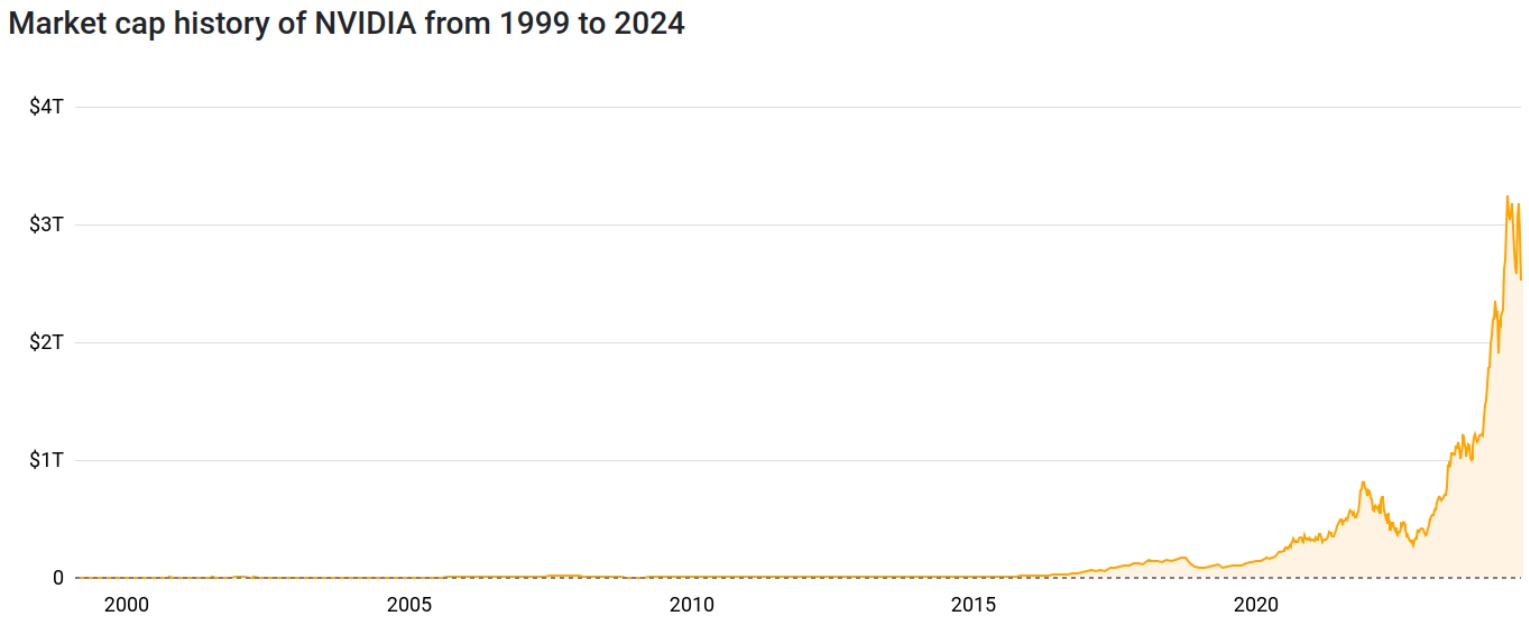

- Independent estimates suggest that "AI"

- requires about half a liter of clean water per query (to cool the server farms)

- requires as much energy per query as it would take to charge a cell phone

- ... hence requiring mothballed coal-fired power plants to be returned to service

- ... and Microsoft recently announced it will pay to bring the Three-Mile Island nuclear power plant – the site of the largest nuclear accident in North America – back into service

- is losing Microsoft $20 per user per month on it's "Copilot" product

- ... and similarly for all major AI companies, so they are all losing tens of billions of dollars (at least) per year

- There's also some interesting geopolitics coming from the fact that the server farms which are needed for "AI" all run, in Europe and the US, on NVIDIA chips (except for Google, which makes its own chips...) – which is why NVIDIA was, briefly, the most valuable company in the history of global capitalism. NVIDIA is based in Taiwan. If China gets more aggressive towards Taiwan, after big tech companies have invested $7 trillion in "AI" based on NVIDIA (Taiwanese!) chips, what happens?

What *is* this "AI"?

The recent technical innovation that made programs producing pretty output which is used to justify all the hype is based on a **Large Language Model [LLM]**.

LLMs consist of a data structure called a "neural network"[1] and two accompanying algorithms:

- one which processes (hundreds of) TB (or more) of *training data* - text and data usually copied from the internet ("scraped") - to create (hundreds of) MB of values in the data structure ("model parameters");
- another which uses the data structure to transform input sequences (the "prompts" people type into a chatbot) into output sequences (which could be words in the same language, or a different language, or pixels of an image) in a way that maximizes the probability of the output given the input and the training data.

[1] in a very poor analogy with biology

- The recent technical innovation that made programs that produce pretty output which is used to justify all the hype is based on something called a Large Language Model [LLM].

- LLMs consist of a data structure called a "neural network", in a very poor analogy with biology. A "data structure" is just a way of keeping information in a computer, like a spreadsheet or dictionary.

- There are also two algorithms in an LLM, one which processes hundreds of terabytes (or more) of training data – text and data usually copied from the internet ("scraped") – to create hundreds of megabytes of values in the data structure ("model parameters");

- and another which uses the data structure to transform input sequences into output sequences in a way that maximizes the probability of the output given the input and the training data.

In other words, a statistical model

Processing sample data to create parameters from which one can calculate output for new inputs is exactly what we call a *statistical model*.

For example, you may have seen a statistical model called a "regression line."

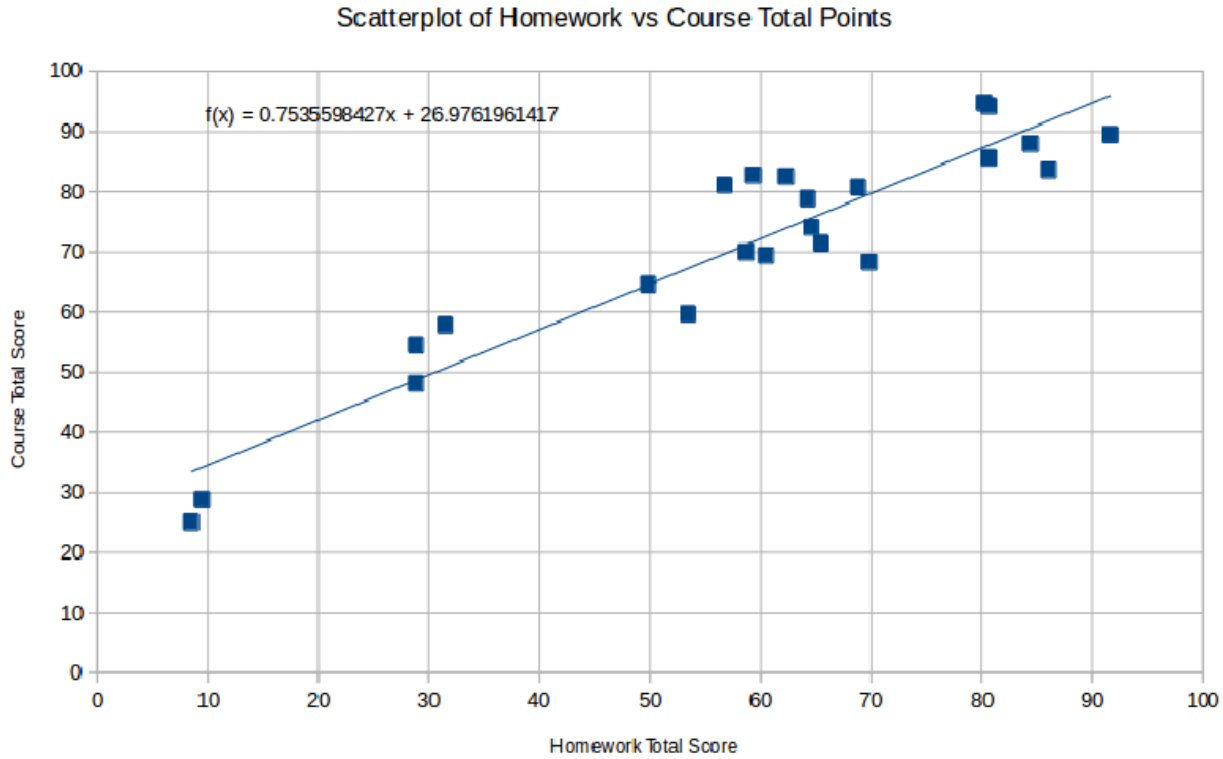

This *scatterplot* comes from an OER I wrote a few years ago. One calculates the coefficients of the equation of the line from the sample data and with this equation one can compute output values ("$y = f(x)$") from new input values ("x").

- Processing sample data to create parameters from which one can calculate output for new inputs is exactly what we call a statistical model.

- For example, you may have seen a statistical model called a "regression line." This one is from a statistics OER I wrote a few years ago. Here, each dot represents a student, and the x-coordinate of the dot (where it is, measured left-to-right) is the student's score on all homework in my class, while the y-coordinate (where it is measured up-and-down) is the student's total points in my class, including exams and projects.

- We can use the line to estimate what would be the most likely total course points for a student whose records were lost, except we happen to know what their homework score was.

- A line is described by two numbers (the slope and y-intercept), so this model has two parameters. And the sample, or training data, here consisted of 44 numbers (x- and y-coordinates for 22 students).

- That's exactly what LLMs do: they don't think, they don't reason, they don't summarize, they don't create .. they just put together the mostly likely output given the sample they were trained on.

- Clearly, if you moved the dots around, the line would be very different. While we don't understand the mathematics of LLMs as well, it is at least clear that the sample they use – their training data – controls what comes out, and if it is biased or racist or doesn't have much text about a particular topic or in a particular language, it will not be able to produce good-looking outputs, no matter what input is given!

Statistical models can do some wonderful things

LLMs are "merely" statistical models... but, actually, I love statistical models: they can do some amazing things. E.g.,

this one

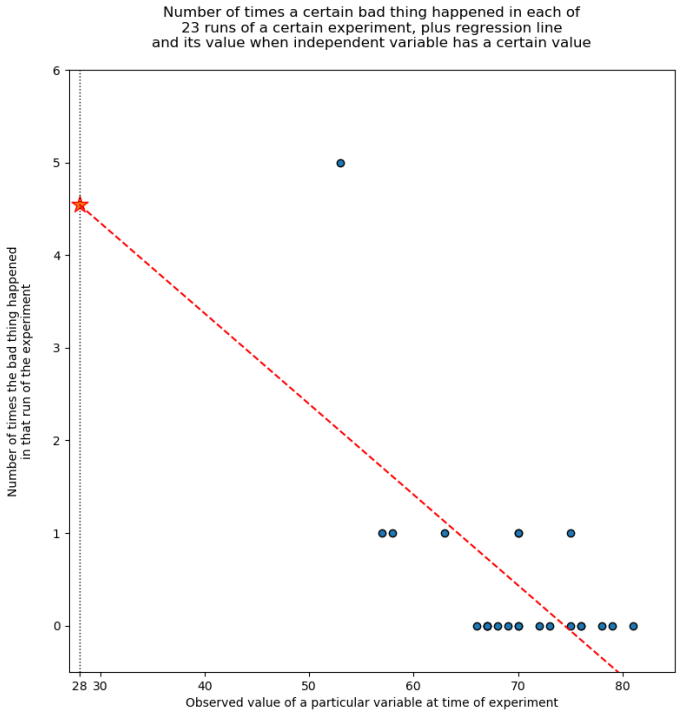

might have prevented

[2]

[2] This image is in the public domain (at least in the United States) because it was created solely by NASA.

- LLMs are "merely" statistical models... but, actually, I love statistical models: they can do some amazing things.

- For example, this is a plot made of the number of small failures in the equipment of the US space shuttle graphed against the temperature at the time of launch. The engineers put this together the night before the Challenger was launched in 1986, and they drew the regression line, and noticed that there would probably be many failures the next day because of the predicted temperature. So they asked that the launch be stopped. But their managers took that slide out of the presentation they did to upper management and NASA officials later in the night, so the decision was made to continue with the launch. The Challenger blew up 73 seconds after takeoff.

Legal aspects: training LLMs

Whether you think the outputs of LLM are great or are mostly terrible (as I do), before we use them in our OER, you might want to think about their legal status – after all, open licensing is one of the main things that distinguishes OER from other (commercial) educational materials. So we have to ask:

When copying and processing the vast quantities of data used to train LLMs, are "AI" companies violating the copyrights of the creators of all of those works?

Some argue that the training of LLMs constitutes merely extracting some facts – the statistics of the sequence of inputs – from works in their training datasets ... and copyright doesn't apply to facts, just expressions. I'm not sure how that holds up when copies of entire works from the training data often show up in LLM output.

Others have argued that training is *fair use* (a doctrine in the US which allows for transformative uses of others' works without their permission, subject to certain restrictions in a famous "four-step test"). But I don't see how training LLMs can be fair if journalists and illustrators are getting fired...

Note: **IAmNotALawyer** and the above does not constitute legal advice! Additionally, my skepticism is not widely shared among legal experts (who, I would like to point out, clearly do not understand how LLMs really work).

- Whether you think the outputs of LLM are great or are mostly terrible (as I do), before we use them in our OER, you might want to think about their legal status – after all, open licensing is one of the main things that distinguishes OER from other (commercial) educational materials. So we have to ask:

- When copying and processing the vast quantities of data used to train LLMs, are "AI" companies violating the copyrights of the creators of all of those works?

- Some argue that the training of LLMs constitutes merely extracting some facts – the statistics of the sequence of inputs – from works in their training datasets ... and copyright doesn't apply to facts, just expressions. I'm not sure how that holds up when copies of entire works from the training data often show up in LLM output.

- In fact, computer scientists have found that sometimes, unpredictably, what comes out of an LLM is almost identical to one of the objects in its training data. So it's not just statistics, sometimes it's the whole thing!

- Others have argued that training is fair use, which is a doctrine in the US which allows for transformative uses of others' works without their permission, subject to certain restrictions in a famous "four-step test". Fair use is what allows one to write a review of a new book in a newspaper that is critical of the book and includes quotes from the book, without asking the permission of the copyright owner.

- But I don't see how training LLMs can be fair if journalists and illustrators are getting fired, that goes against one of the four steps of fair use!

- Note: I Am Not A Lawyer and the above does not constitute legal advice! Additionally, my skepticism is not widely shared among legal experts (who, I would like to point out, clearly do not understand how LLMs really work).

Legal aspects: LLM outputs

Copyright springs into existence (in the US) when someone creates an "original work of authorship fixed in a tangible medium of expression." But who is the author and is it original enough when the work in question is the output of an LLM?
The US Copyright Office has decided this in one case, which, it is thought, may be a strong precedent. Not a great precedent, IMHO.

The reasoning of Copyright Office for the graphic novel **Zarya of the Dawn**...

[cover of Zarya of the Dawn, created by Kris Kashtanova] with Midjourney; in the public domain, at least in the USA, for now.]

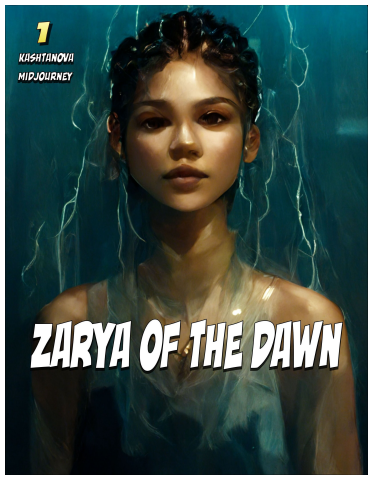

...was based on that of the famous **Monkey Selfie**

[Self-portrait by the depicted Macaca nigra female via Wikimedia Commons, in the public domain, at least in USA, at least for now.]

in that the Copyright Office found the locus of creativity in the LLM's software and said, based on the Monkey Selfie case, that non-human entities, no matter how creative, could not own copyrights.

From my perspective on LLMs as statistical models, this is absurd, and the creativity surely lies in a combination of the creators of the training data and in the LLM user's creativity in using this tool to make a remix. ...But, again, **IAmNotALawyer**.

- Copyright springs into existence (in the US) when someone creates an "original work of authorship fixed in a tangible medium of expression." But who is the author and is it original enough when the work in question is the output of an LLM?

- The US Copyright Office has decided this in one case, which, it is thought, may be a strong precedent. Not a great precedent, in my opinion.

- What happened was that a creator of a graphic novel called "Zarya of the Dawn" applied for a copyright on their work, for which they had written all the text but had gotten all of the images from an "AI" image generator. The US Copyright Office said there was creativity in the pictures, but it was being made by the LLM, not the human, and since there is no provision in the law for a non-human entity to own copyrights, they said the images were instead immediately in the public domain.

- They made reference to a fun case called the Monkey Selfie, where a nature photographer left a camera near some monkeys, one of whom took a selfie. The photographer applied for a copyright in the name of the monkey, which he was going to administer on behalf of the monkey ... the US Copyright Office denied it because of that reasoning about non-human entities and copyrights.

- From my perspective on LLMs as statistical models, this is absurd, and the creativity surely lies in a combination of the creators of the training data and in the LLM user's creativity in using this tool to make a remix. ...But, again, I Am Not A Lawyer.

Legal aspects: legislative actions

A very strong law about "AI" was just vetoed in the US state of California, because the tech industry lobbyists put pressure on the Governor of the state. So that's bad news. The European Union passed something recently which is called the "EU AI Act," which has fairly strong

- protections of the rights of creators (to opt out of their work being used in LLM training corpora, for example),
- preventative measures for draconian uses of LLMs (e.g., no use by police or in schools for things like facial recognition, which is known to be very highly racially biased[3])
- transparency measures about their training corpora, model parameters, etc.

And yet... the EU AI Act defines "AI" as software which shows "nearly human levels of agency and creativity" – so, from my perspective, the Act applies to exactly nothing (today, at least) in the real world, outside a tech company's advertisements!

Although maybe the EU parliament is playing 4-dimensional chess, to get big tech to admit its PR is nonsense or else suffer under the restrictions of the AI Act....

[3] E.g., see **Algorithms of Oppression** by Safia Noble, **Race After Technology** by Ruha Benjamin, and **Unmasking AI** by Joy Buolamwini

- A very strong law about "AI" was just vetoed in the US state of California, because the tech industry lobbyists put pressure on the governor of the state. So that's bad news.

- The European Union passed something recently which is called the "EU AI Act," which has fairly strong protections of the rights of creators (to opt out of their work being used in LLM training corpora, for example).

- The EU AI Act also has preventative measures for draconian uses of LLMs (e.g., no use by police or in schools for things like facial recognition, which is known to be very highly racially biased)

- It also has strong transparency measures about their training data, model parameters, etc.

- And yet... the EU AI Act defines "AI" as software which shows "nearly human levels of agency and creativity" – so, from my perspective, the Act applies to exactly nothing (today, at least) in the real world, outside a tech company's advertisements!

- Although maybe the EU parliament is playing 4-dimensional chess, to get big tech to admit its advertisements are lies or else suffer under the restrictions of the AI Act....

Finding a balance

I have seen very few good uses of "AI" in education ... but not none! Liza Long, in Idaho, USA, is using "AI" in very interesting ways for her writing classes. Alegrìa Ribadeneira, in Colorado, USA, is using it powerfully in her Spanish language classes.[4]

Neither of them just hands the work to be done over to "AI," instead their work very hard and also use the "AI" as a tool. It is not a labor-saving device for them, but a tool which they use along with very thoughtful, human pedagogy.

But does this benefit balance the damage "AI" is doing? For example:

The climate damage is a problem for me.

The theft of other humans' creativity is also a problem.

The abandonment of the greatest academic norm of acknowledging one's sources is deeply disturbing.

A lot of the hype talks as if the LLMs are as good at doing thoughtful, creative work as any human expert – they are not! – which seems to me to be deeply corrosive to foundational structures of truth-seeking and -verification which the world desperately needs at this moment in history.

Moreover, as we have seen, the legal situation is very uncertain and could change quickly.

[4] I'm not sure I know of any other good examples! You can find their work by searching on the Internet.

- I have seen very few good uses of "AI" in education ... but not none! Liza Long, in Idaho, USA, is using "AI" in very interesting ways for her writing classes. Alegrìa Ribadeneira, in Colorado, USA, is using it powerfully in her Spanish language classes. ...I'm not sure I know of any other good examples! You can find their work by searching on the Internet.

- Neither of them just hands the work to be done over to "AI," instead their work very hard and also use the "AI" as a tool. It is not a labor-saving device for them, but a tool which they use along with very thoughtful, human pedagogy.

- But does this benefit balance the damage "AI" is doing? For example:

- The climate damage is a problem for me.

- The theft of other humans' creativity is also a problem.

- The abandonment of the greatest academic norm of acknowledging one's sources is deeply disturbing.

- A lot of the hype talks as if the LLMs are as good at doing thoughtful, creative work as any human expert – they are not! – which seems to me to be deeply corrosive to foundational structures of truth-seeking and -verification which the world desperately needs at this moment in history.

- Moreover, as we have seen, the legal situation is very uncertain and could change quickly.

Contact Info

Email: jonathan@poritz.net

Fediverse: @poritzj@mastodon.social

These slides:

poritz.net/j/share/OLaGAI4OE.pdf

and all files for remix:

poritz.net/j/share/OLaGAI4OE/

Or else use ⟶

to go to my homepage, where you can follow links to all of my shared materials.

- Thank you for your attention.
- Here is how to contact me, I'm always happy to discuss individually anything that might be useful to you!
- Thank you so much to the interpreter!