AI and the Commons: Solutions to be found only beyond licences

By John Weitzmann, Legal Counsel at Wikimedia Deutschland

Prompted by the EU being in the midst of regulating Artificial Intelligence as a field of technology, the blogpost to which this is a reply, calls for reassessment of what Open Content can be used for and of whether we actually want such uses to happen unchecked around AI.

It identifies uses like the training of facial recognition systems as harmful and, in a way, thus contrary to the idea of sharing as a way of caring. And it attributes the possibility of Open Content being used for such partly harmful means to the reduced level of copyright enforceability brought about by open licensing, i. e. by the famous "some rights reserved", which leaves less leverage to forbid harmful uses.

While we surely must not shy away from looking at what develops with and around Open Content, and from looking for solutions for harmful effects, this reply argues that such solutions cannot be found on the licensing level, and must not try to leverage copyright as a prohibitive means – unless we are willing to sacrifice the idea of Open Content altogether.

New technologies means new dark sides

Many uses of facial recognition systems live on the rather problematic end of the scale that reaches from usually benign at one end to potentially harmful at the other. A whole new level of surveillance is made possible by automated identification of people in the live footage of an ever increasing number of cameras, increasing both in public and private spaces. This must be a concern for all those who advocate for strong fundamental rights, and among these are the advocates of Open Content and open content projects.

One might argue that facial recognition is not actually an application of artificial intelligence technology, but is rather a sophisticated method of pattern recognition and an instance of deep learning mechanisms. But that is besides the point. It often enough will be part of larger technological systems that involve AI, providing information to them, and in any case is in itself a technology with an enormous potential for abuse.

We should maybe clarify as well as widen the scope and look at: Digital content used for enhancing autonomous systems or automation in general (the term Automated

Decision-Making, ADM, comes to mind). The breathtaking potential of automated systems includes a likewise breathtaking potential for abuse.

Nobody with stakes in digital technology, the internet, and fundamental rights must therefore disengage from the debate around such systems and how to regulate them.

At the same time we have to be quite precise about what type(s) of content we talk about. That is not only because the property of it being open (in the meaning of the open definition and freedomdefined.org) is key here. Also the possible means for regulation differ regarding the content in question.

Openness doesn't mean to leverage, but to lessen control

An openness that says "here's content that you can use, unless your use is something we don't approve of", is a very peculiar kind of openness, to put it mildly. That is not to say that the reasons for disapproval are always wrong or harmful in themselves. The example of training use for facial recognition is actually a well-argued one, even though facial recognition can of course also be used for many beneficial applications.

Let's take weapons instead: Even here some might claim beneficial effects, but a very large majority would probably deem the use of content in the course of weapons development something (literally) harmful. But to limit the usability of content regarding such uses makes "open" into "open for certain uses". One doesn't need to look at the friction produced by the infamous non-commercial condition present in some of the CC licences to understand this, but it is a very telling example:

Even though NC can make sense as a means to restrict certain uses / users, it mostly hits the wrong ones. In other words: If you narrow your view on a tiny group of potential users you want to control and deflect, restrictions can make a lot of sense. And that, by the way, is the idea behind the classic IP licensing business where exclusive rights by default exclude everybody except a tiny group that bought a licence.

But if one takes into account all potential users, which is the perspective the paradigm behind Open Content usually takes, such restrictions limit a number of users far greater than intended. Depending on the context, the ratio is probably that of 2% reasonably restricted vs. 98% collateral effect, although no conclusive research seems to exist on this point.

Some perceive the open licensing approach as something that already makes a very cunning use of the infamously strict copyright mechanics, something that is turning the tables on those despicable copyright control freaks. Maybe that's one reason for being overly ready and willing to tinker around with the licence grants as leverage? But when looking closely at the 'licensing' in open licensing, it becomes apparent that the approach contains hardly anything that justifies the word "leverage":

What free / libre / open licences mostly just require as a condition is attribution, as laid down for example in the BY clause of Creative Commons licences. Yet, attribution is arguably one of the core rights in any copyright codification, one that thus would be required by default anyways, by law. Clauses like in CC BY shape that requirement by making it more explicit and detailed, and might even in some cases make attribution better enforceable, but the difference is really not that big to what would be required by law already. There's definitely not much 'leveraging copyright' here.

But there is a second condition possible in licences that are meant to qualify as free / libre / open: Copyleft (or as it is known in Creative Commons licences: Share Alike). Just as the attribution condition, copyleft only requires a certain positive action to happen, on top of the uses that are allowed in the licence grant, thus it too is far from prohibiting anything or limiting the licence grant in any relevant aspect. Still, the positive action it requires goes much further than what the law requires anyways for example regarding attribution, see above. It requires that content derived from the initial Open Content to also be openly licensed (if that derivative content is in fact published, that is). So yes, there is some leveraging in copyleft, helping to preserve the openness through stages of derivation. But in no way does this limit what kind of uses the licensee is allowed to engage in.

And, this comparatively mild positive action leverage alone is already producing considerable problems: Many people, even lawyers at times, have a hard time understanding what copyleft requires of them exactly and when it is triggered. This is a strong indicator as to just how susceptible open licensing is to complicating factors. And this one only relates to the level of positive action. The more drastic example is of course the already mentioned non-commercial clause, which relates to the licence grant itself. It makes a licence that contains it non-open altogether. Similar effects loom whenever limits are put on the "whatever" in the licence grant, because they all by principle require interpretation and produce grey areas of legal uncertainty.

How big of a chunk is it anyway?

Being in the middle of arguing about the use of Open Content for training automated systems through deep learning, one very basic question cannot be avoided: How relevant really is the chunk of training content that comes from the world of Open Content?

We can't directly know the answer, nobody can. But we can make an estimate, because there's little doubt that people responsible for getting trainable systems in contact with content to train on will simply go for the maximum amount of content they can get. Thus, the ratio between Open Content on the one hand and all internet content that is "freely" (as in unhindered) available on the internet on the other will most probably be indicative of how much training content is actually Open Content.

Creative Commons regularly published estimates on the overall size of the commons, meaning on the overall number of CC-licenced works present online. Those figures are impressive for sure: There don't seem to be up-to-date figures in CC's "State of the Commons" report format, but even the latest numbers from 2017 range in the lower one-digit Billions, with an impressive growth rate. Yet, looking for example at the picture content that is so very important in general and for the topic at hand in this text, compared to the overall amount of pictures on the internet, including social media platforms and such, the amount of CC-licensed works is tiny.

And, as mentioned above, we must assume that AI systems in general and facial recognition in particular is being trained on whatever can be scraped off the net. Google's mothership company Alphabet can train its own systems on the entire indexing database of Google Image Search (plus probably a lot more coming from Google services withing the Android mobile phone ecosystem). "Can" in this context means "being able to" technically, not: "being allowed to". And let's be honest: Being able is the limiting factor here, it's very unlikely that a relevant amount of imagery is sorted out due to legal prohibitions – first and foremost because hardly ever can anyone outside any given company that is interested in systems training see or otherwise know what happens inside, what content databases are available to them and how they use them.

A Commons is not about ownership

Some might say: "But still, we must at least take care of what is done with our content, even if that's only a small chunk!". Well, here's the hard truth: It's not our content. That's the thing about Open Content that is both wonderful and annoying, that it belongs to all and thus to nobody in particular. It was given not to us, but to the commons – each content piece individually by individual initial rights holders.

There is no central hub or institution here. The commons is a massive amorphous cloud-like thing and no licensing stewardship org (like Creative Commons Inc, the Free Software Foundation or the Wikimedia Foundation) nor any group of open licensing enthusiasts or activists is owning the commons, or can even at least claim that they are tasked with custodianship for the commons or with taking care of it.

The late Silke Helfrich was a strong proponent of the idea of commons caretaking, of a kind of custodianship for the commons. But if such custodianship exists, it rests with precisely those who have contributed to the commons, meaning, if the global commons in question in this blogpost consists of creative works contributed by one million people, that very group of a million heads is its distributed custodian. This multi-headed custodian isn't in any position to form a combined will, not even by majority vote, let alone unanimously.

All these people contributed to the commons under a notion that says: Open means that everybody can use this content for whatever they please. Note that built into this is of course the effect of other rules that do exist outside copyright, which means that certain uses might be

prohibited according to such other rules – but not according to intellectual property rules a.k.a. Copyright. Yes, some conditions apply, but they do in no way limit the rights grant. They instead produce positive obligations on the user's part to do something.

The rights grant is all-encompassing, when it comes to truly open / truly free licences. The counter example is the CC NC clause, which does limit the rights grant and thus removes the licences that contain it from the world of free licences. And, this all-encompassing notion of Open has been unaltered for decades. That's why the entire amount of content contributed since, regarding software, the 1980s and regarding other works since around the early 2000s rests on this all-encompassing notion of Open. To move it to another notion of Open, one of "use it for whatever you want except A, B and C" isn't possible. There's no way to gather the entire group of contributors and decide on this.

In other words: To change the notion of Open underlying the existing digital commons of Open Content retroactively, is as impossible as removing the foundations of a building after it was built, and substituting them with something else. In the case of an immaterial commons defined by rights grants, one might even argue that it's not only the foundation that would be substituted, but all the bricks as well.

Coming from this building metaphor, it should be quite easy to see that changing the notion of 'Open' in Open Content would mean nothing less than starting a whole new digital commons and departing from the old, like a gigantic fork of code, a schism of Openness. Given sufficient reason for such a schism, it shouldn't be deemed an unthinkable thing to do. The question is, whether countering undesirable effects of facial recognition technology, that is indirectly supported by the existence of Open Content, is sufficient a reason. The author of these lines thinks it's not, in particular because there in fact are other means outside the licences, which will be mentioned at the end of the text.

A mandate on the input side, sure, but hardly any on the output side

What's more: Due to Open Content decisively not being "ours", there's a huge question mark regarding the reach of our political mandate when we engage in advocacy. Sure, on the input side there is something could be called a mandate, one coming from the very idea behind Open Content: Built into this idea is the view that more of the open stuff is better than less, which can be understood as giving us – or anyone else, for that matter – the mandate to strive for expanding the Open Content world.

On this input side of things, there are several more specialised mandates thinkable, one for example regarding Free Knowledge (Hello Wikimedia folks!), one for enlarging the larger digital commons of cultural artefacts (Hello entire group of openness activists at large!).

On the output side, however, when it is about what to do or not do with the ever growing digital commons things are different. There is a kind of unsurmountable wall between the input side

and the output side, and that is the conscious decision, made by whoever contributes, to hand something over to the commons. This decision determines what the scope of permission is for later. It thereby also determines which of several possible commons the content in question becomes part of. And it cuts us (and everyone else) off from changing the scope of permission later, on the output side.

It follows that nobody has a mandate to limit or widen the scope of permission that was chosen by the contributing person at the point when a piece of content joined the commons. And, as was elaborated on above, the permission scope of Open Content as we know it has no limits as to what to do or not to do with the content.

It's thus highly questionable whether it is the job of Wikimedians to purposefully engage in political framing outside of the core purpose of the Wikimedia projects. Working towards "*... a world in which every single human being can freely share in the sum of all knowledge*" (from the Wikimedia Foundation's motto, see their website) does definitely not extend to influencing what all those humans do with that sum of all knowledge. It cannot extend to there, and we must not pretend that it ever could. And to regulate, or have regulated by the EU or any other body, what others can or cannot do with knowledge means going beyond the freely sharing idea and our mandate.

Let's not run into severe contradictions for chasing a ghost

But let us for a moment pretend, things were different in several ways: Let's pretend that Open Content were a rather significant chunk of the training material for automated systems. And let's further pretend that we could reasonably call that Open Content "ours". And let's also ignore the fact that open licensing mechanics are so easily ruined by detailed restrictions and the interpretation and legal uncertainty they bring.

Even then we would need to ask ourselves, whether the copyright leverage we'd like to install in order to push back against the use of Open Content for training facial recognition systems would actually work in practice. I would argue that it wouldn't, simply because it hardly ever comes to light which pictures (or other content pieces) were used for which training of AI by whom and when.

The uses in question here happen out of sight, in an unknown number of blackboxes that are the IT systems of all those partaking in the race for ever more powerful automated systems. Nobody who refrained from making their digital content technically hard to scrape from their website can ever be sure that their content was not used in training automated systems.

So, do we count on leaks? Conscientious employees of IT companies large and small who step forward and reveal "There is content in our training database that we don't have a clear permission to use for training!" is an unrealistic scenario. If that is ever going to happen, it would represent single grains of sand on a beach.

Plus, to want and encourage people to come forward in that way would have us run into a severe contradiction, because it is not at all clear that what we would want to have regulated here, i. e. the drawing of technical conclusions from analysing content, is within the field of what can be regulated by intellectual property rules. Quite to the contrary we have for many years claimed the opposite.

For almost two decades now we strongly advocate the claim that “The right to read is the right to mine”. By this we mean that there is no legal anchor to forbid or allow anyone to look at content and draw their conclusions from doing so. We put this forward not only in support of science made directly by humans reading stuff, but very prominently also regarding Text and Data Mining. In other words, we told the lawmakers far and near that they cannot, or at least must not try to, regulate to the effect that it can be forbidden to use technical systems to simply “read” and analyse content on a massive scale. Well, training computer systems on data is the very same thing. These systems take in content and deep learn from it.

And there’s more, because we also, and probably even more forcefully than for TDM, advocate for not using copyright as a means to achieve external means that are not part of copyright’s intended scope. Our classic foe here is de-facto censorship, in Germany we were even successful in coining a framing term for this: Zensurheberrecht.

In trying to achieve limiting effects around AI abuse, we are ourselves moving into leveraging things outside of copyright. Just as copyright is not intended to regulate what a society can look at, meaning censorship, it is not intended to regulate what machines can look at, meaning automated systems training. If we would use the licensing grant as leverage here, we would engage in the very same practice we so far so passionately despised – just because now we suddenly see an external regulation goal we deem beneficial.

Apart from the problem with an insufficient mandate, see above, to go down the path of legally restricting the ‘Open’ in ‘Open Content’ would be nothing less than a slippery slope scenario and a floodgate argument mixed together into one of the most toxic remixes thinkable in our part of the advocacy world. Certain content industries are only waiting for such an open flank. Let’s not give them that.

Not just a nerd’s delight: Licence proliferation

And finally, there’s one rather legalistic topic that is often pushed aside as purportedly only a nerdy one: Licence proliferation. But it matters as little or as much as one wants to use licence texts as a means of crafting control mechanisms. So, anyone who plans to tinker around with the rights grant as the core piece of Open Content licensing needs to face this thing.

Licence proliferation in the context of Open Content means a situation where the overall pool of content that is freed up via licences and that is at least potentially meant to be remixed or

otherwise used in combination, is licensed under an increasing number of different licences. This hinders the functioning of Open Content ecosystems due to two effects of the proliferation: Increased complexity and licence incompatibility.

Important background to understand the incompatibility part is the fact that any substantial difference between two licences, i. e. differences concerning either the scope of the rights grants or the conditions under which the rights are granted, tends to make these licences protective against each other. Incompatibility in this sense means that at least one of the licences cannot be adhered to without breaching the terms of the other – sometimes this even goes both ways. There is a kind of dirty hack to fix this, which consists of explicitly declaring, in the licence text itself, the licences compatible with each other. Such formal compatibility doesn't go too far in fixing the problem, though. It rather covers it. And it doesn't help at all with the complexity issue:

Open Content is all about enabling, about enabling even non-lawyers to freely engage in things like creative expression, without the barbed-wire fences prohibitive copyright laws erect everywhere. The idea of Open Content is to lower the so-called transaction costs to a level where even non-profit projects and individuals can dare to handle copyrighted material in a reasonably safe way. That requires simplicity. The fact alone that incoming content is governed by more than one licence, i. e. rule set, can increase the transaction costs to a level that is prohibitive, especially if the protagonists of the project are good people who don't want to break the law.

Prohibitive means in such cases that the projects simply don't happen. We've seen this numerous times for example with smaller educational projects, where only a handful of (possibly unjustified) infringement claims brought down the whole thing. That's what complexity does to Open Content projects, and complexity is what comes from licence proliferation. With those who are less pressed to be copyright compliant, if the licence landscape becomes too complex, adherence to licence conditions becomes arbitrary and eventually random. The whole approach deteriorates and eventually becomes obsolete.

So, from the perspective of those who want to rely on actual control to reduce uses they consider harmful (like using Open Content as training material for facial recognition), licence proliferation must be a major concern. Licence proliferation itself is entirely agnostic regarding the content of licences. In other words: Licence proliferation happens regardless of what is changed in the licence's texts. Even the "versioning" that Creative Commons did between their licence set 1.0 and today's 4.0 in essence proliferates, which is why so much effort was put into version 4, so that it could hopefully be the ultimate final set.

Historic fixtures

Given this agnosticism it would be completely fine – from a licence proliferation viewpoint – if there would be, say, only two widely accepted standard licences that would inter alia prohibit to

use the licensed material for facial recognition training, or for use in weapons systems or any other use deemed harmful.

But that is not the case and nobody can turn back time. The open licences that we have, with CC BY and BY-SA as the de-facto standard for general IP-protected types of content, as well as the data/database-related licences developed by others than CC and the FOSS licences, do not feature such restrictions. They do so for the reasons given above, but even those reasons aside, it would require a complete retroactive re-set of the Open Content world to introduce anti-AI-training prohibitions without destructive licence proliferation. This is factually impossible. And it would also require a complete re-licensing of existing content, lest we lose that content for the future. That is undesirable.

Some have made suggestions to the effect that CC should introduce added restrictions in the course of versioning to a Version 4.X or 5 of the suite. But added complexity in the form of yet another version aside, that doesn't work, as explained above: Any additional restriction on the material level of the rights granted destroys compatibility vis-a-vis the previous versions. So, unless someone manages to travel back 20 something years and introduce specific restrictions into CC licenses back then, or even better, travel back around 40 years and introduce them into the FOSS licenses at the beginning of the emergence of Open Content, we are stuck with the very liberal rights grants we have, at least in the Open Content mainstream.

To sum up: Even on the legalistic level a move to limit the 'Open' would constitute nothing short of a break in the entire Open Content world, with an incompatible before and after. While admitting that there are harmful uses that indeed are, albeit in a rather minor way, facilitated by Open Content, the author of these lines thinks it a good thing that we are stuck with this very wide spectrum of allowed uses.

Instead of breaking the Open Content approach while trying, we'll have to deal with harmful uses through other regulatory means, and not only with regulation, but also on the stage of social norms and in the arena of societal discourse. As is always the case with freedom, because freedom can always be abused.

So, should we refuse to engage with the topic of AI and The Commons? Not at all. We should, however, refuse to instantly jump on copyright-based solutions, for they easily destroy the Open Content mechanics at their very heart. Instead we must look for solutions beyond the core rights grant of open licensing, for solutions that are much closer to the problem. And yes, they do most probably exist:

If we are discussing imagery that can be used to deep learn about faces, that is less about those who took the photos, so less about copyright holders. It is much more about the people depicted on the photos. If we want to leverage law to go against such uses, it stands to reason not to look at copyright, but at personality rights, data protection and privacy law in general. Open Content keeps out of these fields of law as much as possible, which is why we are much

less likely to inadvertently damage Open Content as a concept when advocating regulation in these more appropriate fields of law. They, by the way, also have a much stronger tradition of collective action than copyright.

Conclusion

Restricting certain purposes of use is alien to the Open Content definitions we have. Unless we want to risk a schism of Open, dividing the Open Content world into two and boosting licence proliferation, we are stuck with the liberal definition Open.

Even if we found running the schism risk worthwhile, we could hardly claim to have a mandate for such a move, a mandate on the output side of Open Content, the side where the uses happen. But even if we wanted to risk the schism and had a mandate for that, we'd be chasing a ghost solution, because neither does Open Content represent an all too relevant chunk of the training material for automated systems, nor would a legal prohibition clause have much practical application and effect, as most violations will never become known.

At the same time we'd flagrantly contradict our own advocacy, which since many years supports the slogan that "the right to read is the right to mine" and fiercely calls for not leveraging copyright for purposes external to it. Instead we should engage in pushing back against problematic developments around digital content – if we think our mission covers that – in ways that do not touch the core rights grants in open licensing. We should focus on personality rights and privacy law. These fields are closer to the facial recognition scenario and feature means that might prove much more effective than tinkering around with masses of individual rights grants.

Berlin, March 2022

John Weitzmann serves as General Counsel of the Wikimedia Chapter in Germany, where he previously also re-built and headed the advocacy team.

For an abbreviated version of this article visit wikimedia.brussels