Briefing on Open Source AI

On 10 November 2023, 17:30-18:30 CET, Women in AI Austria hosted a panel discussion with Zuzanna Warso (Open Future) and Alexander Baratsits (Creative Commons Austria, cba.media), moderated by Katja Mayer (University of Vienna), to go deeper into the challenges and opportunities of open source AI.[1] The discussion centred on the three themes of current regulatory developments around open source AI, the challenges around open washing, and finally, some best practices and positive outlooks for this developing field.

In this briefing, we will sketch some of the central themes of open source AI systems and attempt to roughly characterise a broad field marked by a variety of debates. Rather than define what open source AI is, we will try to show how open source AI relates to particular kinds of issues, concerns, movements and values. This briefing has been updated after the event with the inputs provided in the course of the discussion.[2]

[1] https://www.linkedin.com/events/policychat-exploringopensourcea7115987928379535362/
On short notice, Jeannette Gorzala (European AI Forum, AI Austria) was unable to join the discussion as initially planned but has contributed to this briefing in writing.
[2] The briefing published prior to the event is available on the Women in AI Austria website: https://www.womeninai.at/uncategorized/background-briefing-on-open-source-ai/

What you will find in this briefing

Discussion and outlook

Our aim in this section is to briefly synthesise the discussions we had, but also to elaborate and build on them. We have therefore chosen to integrate direct quotes from the discussion: Zuzanna's statements are formatted *in italics and green* whereas Alexander's are formatted *in italics and blue*. Jeannette included some comments on the discussion as well, formatted *in italics and purple.*

Regulation

Most discussions around regulation in Europe currently veer towards the AI Act – and open source AI has become a central topic of debate to this proposal, still to be finalised as of writing.[3] Amongst other things, undifferentiated understandings of openness meet in this debate with a conflation of foundation models and open source, which is why we believe these issues merit some teasing apart and sorting out – especially because open source AI systems are set to be exempt from some requirements of the AI Act and it remains unclear whether this refers to models where the source code or only the model weights have been published, which would limit scrutiny and community governance.[4]

On the one hand, openness leads to more decentralisation in that the control over a technology is more dispersed than in a proprietary context. Decentralisation and governance do not always go easily together, and in the case of the product safety approach chosen by the AI Act, decentralisation – as we have seen it in open source initiatives – undermines the categories of producers and users (and may well raise issues with the AI Liability Directive further down the road). But decentralisation is also an opportunity. It allows more people to contribute to AI systems development, scrutinise models and applications, and participate in deciding which AI systems should be developed and how. (By the way, this is why Women in AI Austria picked up the topic: we believe a more open community offers more pathways to diversity and more opportunities for equitable development!)

However, the labour and care work that goes into making and maintaining open source products is susceptible to exploitation. This is what Zuzanna described as the *paradox of open*,[5] which is to say that *openness can both challenge but also enable the concentration of power*. Openness can easily be abused by companies, as when the content produced within the digital commons is then used to train proprietary AI systems. There is a wider question to be asked here about how to protect openness – how to ensure, in effect, that the benefits of shared voluntary labour accrue to all and not to one or some –, but it does show that *openness is not a magic bullet to the heart of power*.

Taking a step back to look at what regulation is trying to solve, we need to keep in mind that the AI Act's first concern lies with risks arising from products in the context of particular use cases. Open

[3] Open Future has published extensively on the intersections of open source and the AI Act, e.g. https://openfuture.eu/blog/code-is-speech-and-speech-is-free/ or https://openfuture.eu/wp-content/uploads/2023/07/230725supporting_OS_in_the_AIAct.pdf.
[4] More about the difference between open weights and open source here: https://promptengineering.org/llm-open-source-vs-open-weights-vs-restricted-weights/
[5] Find out more in Open Future's take on the paradox of open: https://openfuture.eu/blog/open-source-ai-and-the-paradox-of-open/

source AI systems are not inherently more or less risky than proprietary AI systems, just like certain AI technologies (e.g. deep learning) are not more or less risky than others (e.g. search algorithms); the risk, as seen from the AI Act, emerges from particular uses of AI systems. And open source AI systems are becoming integral to civil society initiatives or commons-focused projects, as Alexander argued. Evidence-based regulation therefore needs to find ways to enable open source AI development and use without losing track of the adverse impacts of certain use cases.

Notwithstanding these issues, openness can do some things for regulation. For instance, transparency and the wide availability of documentation are cornerstones of both the AI Act and open source initiatives.[6] Enabling open source AI development and use could also be a regulatory intervention for a different set of purposes, namely to counteract the concentration of power. But not only in terms of avoiding power concentration: *openness enables regulators and market participants to learn more about the specific potential hazards of foundation models, respectively general purpose AI, enabling better and more targeted measures.* In order to do that, however, a different set of problems needs to be addressed, namely *access to compute, data and development resources*. This shows that on the other side of regulation via the AI Act, larger questions of infrastructures for AI development and use need to be resolved, for instance through public funding for open source endeavours. The two main bottlenecks of data and compute, for instance, *require infrastructure and also maybe new ways of governing these infrastructures*, also to make sure that science can remain independent from the private sector. *The case of OpenAI is an example, where initial open source developments 'go private' due to funding gaps, ending in an unfortunate public display of the conflict between commercial and non-commercial company goals.*

Of course, to which extent the AI Act will apply to open source AI systems depends significantly on what is, in the end, defined as open source (or even AI systems, for that matter). As these processes of definition by the open source community and the European legislators are running in parallel, much of the effect of current attempts at regulation will only become clear once these moving parts have settled.

Open washing and the challenges of balancing openness and regulation

Some days, it seems like the field of AI is very partial to laundry – and so we can add to our concerns about ethics-washing the phenomenon of open washing. Zuzanna described this as the practice of claiming openness for the sake of good publicity, popular support etc. whilst undermining or continuing to undermine the very principles of openness. This is the case with a number of purportedly open AI systems, like Llama 2, which in fact have very restrictive licensing that *includes anti-competitive measures that basically break the established norms around open source and do not provide the freedoms that open source is supposed to secure*.[7] One of these norms that open washing breaks with is access to documentation and transparency. AI systems consist of several components: the model which includes the parameters and weights, the software which runs the

[6] Notwithstanding these principles, there remains a debate about what transparency actually is or should be, and of course how it could be measured. On Stanford's Foundation Model Transparency Index, for instance, see EleutherAI's critique: https://blog.eleuther.ai/fmti-critique/

[7] See Open Future's detailed analysis here: https://openfuture.eu/blog/the-mirage-of-open-source-ai-analyzing-metas-llama-2-release-strategy/

model, and in the case of machine learning, the data used to train the model. In order for a community to work with the open source AI system, these components need to be both available and documented. In practice, especially information about training data is lacking, which inhibits use of the AI system, because *meaningful accessibility requires information about data*.

An AI system that is released as 'open source' is not necessarily open with regard to all system and risk components. There are different gradients of openness, e.g. related to the openness of model weights versus the openness of the source code. In this context, it is crucial to note that openness may be a way of avoiding regulation and that the open ecosystem increasingly includes companies which stand to gain from different forms of enclosing value.[8] These factors add to the complexity of the phenomenon of open-washing, calling for increased attention to what is termed open by whom and what these levels of openness imply for governance.

Against this backdrop, where does that leave open content distribution providers? On the one hand, Alexander continues to recommend CC-BY licenses to those seeking to distribute content, certainly from an open access and open knowledge viewpoint but also because it allows this content to be easily picked up by other initiatives such as Wikipedia. However, licenses for data open up questions around AI applications – when data have a non-derivative license, for instance, can they actually be used for training automatic speech recognition or translation?[9] Open washing gains a different dimension here: openness for data use may not correspond to openness for AI systems, as the openness of code differs significantly from the openness of data, both in practical application and regulatory frameworks.

Moreover, there are wider questions of value attribution at the intersection of data and AI systems. Openly accessible data can be integral to developing AI systems,[10] but value created through the AI systems usually does not make it back to the open ecosystem on which the AI system was built. For this reason, there *should still be this open access approach, but it would be nevertheless fair if there would be some rewarding system in place*. For Zuzanna, the matter of *giving back to the commons and taking care of the commons are some of the key principles that should be implemented*, although there are plenty of open questions about how to do that in practice (pun unavoidable). When we evaluate open washing and hold it up against openness, we need to remain mindful of the effects of particular uses on the open ecosystem as a whole: it is crucial to consider the impact of specific practices on the overall open ecosystem and the diverse interests embedded within it.

[8] See for instance the newly founded AI Alliance: https://newsroom.ibm.com/AI-Alliance-Launches-as-an-International-Community-of-Leading-Technology-Developers,-Researchers,-and-Adopters-Collaborating-Together-to-Advance-Open,-Safe,-Responsible-AI

[9] These questions arose in the context of the media data space which cba.media is trying to build: https://cba.fro.at/building-a-european-cultural-backbone

[10] Not all AI systems require large amounts of data for training; when thinking of AI systems used in industrial contexts, for instance, these would not necessarily require open datasets, although e.g. weather data is often integrated into a broad range of applications.

Bright sides and outlooks

Putting the indeterminacy of regulation and the complexities of openness aside, there are quite a few practical opportunities for developing open source AI systems in an open ecosystem.[11] One of these lies in using AI systems to make open content more easily discoverable and usable, which is one of the use cases of the Fair By Design project[12] that Alexander is working with. The aim of this project is to develop a safe and unbiased recommender system by building in checks, which necessitates close engagement with the stakeholders. For us, this points to the continued need to engage with questions of participation: who gets to participate in developing technology and how? In this context, *open source AI has had some benefits that other ways of developing technologies do not because open source is also about communities cooperating with each other, taking each other's work, also being accountable to each other*. And next to sharing technology, this also means sharing information about governance experiences, as Hugging Face has done,[13] but also making sure that people know that technology is accessible and not *hocus pocus*, as some might like to believe.

The project Display Europe,[14] on the other hand, seeks to develop a media platform for European content in 15 different languages. Here, the backbone builds on the metadata provided by different European NGOs to form a federated portal for their content. In turn, this initiative forms part of the Trusted European Media Dataspace (TEMS),[15] a collaboration between European media companies and also commons-oriented initiatives like cba.media. One of the contributions of Display Europe here is to ensure interoperability between the data models.[16] And within this context, one of the unresolved questions relates to remuneration: with the origin data not on equal economic footing, but now connected through a shared data space as well as search and recommendation based on AI systems, how do you ensure that the value generated by the collaboration is equitably distributed between commercial and commons-focused partners?

Just like open source and open data initiatives have formed a valuable part of our information ecosystem and common toolbox, open source AI should hold a place that is both sustainable for those collaborating on its development and use and inclusive to allow for a wide range of participants.

Concluding Reflections: Exploring the Landscape of Open Source AI and Regulation

Concluding, we provide reflective insights and a broader perspective on the complex interplay between open source AI and its regulatory environment. By revisiting the key themes explored in

[11] Courtesy of Open Future, an overview of applications for open source AI systems: https://docs.google.com/spreadsheets/d/1kt5jp1U50AfDGEKsqFXvtCCQ8HYvQGcya1vf_OaD8x0/edit#gid=0
[12] More details here: https://www.fairbydesign.eu/the-consortium
[13] Jernite, Y., Nguyen, H., Biderman, S., Rogers, A., Masoud, M., Danchev, V., Tan, S., Luccioni, A. S., Subramani, N., Johnson, I., Dupont, G., Dodge, J., Lo, K., Talat, Z., Radev, D., Gokaslan, A., Nikpoor, S., Henderson, P., Bommasani, R., & Mitchell, M. (2022). Data Governance in the Age of Large-Scale Data-Driven Language Technology. 2022 ACM Conference on Fairness, Accountability, and Transparency, 2206–2222. https://doi.org/10.1145/3531146.3534637
[14] More details here: https://displayeurope.eu
[15] More details here: https://tems-dataspace.eu/
[16] How to govern data spaces in an open manner is another question on which Open Future has shared some reflections: https://openfuture.eu/note/talk-on-commons-based-ai-dataset-governance-for-osi/

the discussion, we hope to shed light on the nuanced challenges and opportunities that define this evolving landscape.

The paradox of openness challenging and enabling power concentration lies in the fact that while it democratises access and contribution, it can also be exploited by dominant entities to consolidate their control, often by leveraging freely available resources for proprietary gains.

The European AI Act is a key regulatory focus, which will have great impact on AI governance globally, however it is important to distinguish between foundation models and open source AI. Decentralisation, a result of openness, currently conflicts with governance structures proposed by the AI Act, affecting categorisations of producers and users. The AI Act emphasises risks from AI products, but it is important to note that open source AI systems are as risky as proprietary ones, recognising that the risk level of open source AI systems is context-dependent. Above all, it is crucial that the AI Act is very specific about its terminology when it comes to openness to ensure that community control will be possible.

“Open washing” involves falsely claiming openness for publicity, often accompanied by restrictive licensing and inadequate transparency in documentation and data.

Panelists agreed that effective regulation must balance AI Act compliance with supporting open source AI development, focusing on improving access to resources and infrastructure. Additionally, they emphasised the need for enhanced methods of value attribution to support open ecosystems, noting that although open data is often crucial for AI development, the generated value frequently fails to reciprocate to the open ecosystem, thereby raising concerns about fairness.

In summary, open source AI presents significant opportunities for developing technology for social good and fostering community-led initiatives, highlighting the need for inclusive participation and fair value distribution. Projects such as Fair By Design and Display Europe not only showcase practical implementations but also underscore the necessity of establishing appropriate governance measures to guarantee equitable distribution of value.

Background

Open source movements

The first association made in connection with open source AI systems usually relates to open source movements, particularly around software. Initially, software was developed in an academic setting, bound to hardware but shared relatively freely within the community. Later, from the 1970s onwards, when computers became more widely available, software was developed into a proprietary product separate from the hardware; through these developments, software became subject of copyright laws, and thus entered the same legal domain as works of art. Pushback against this course of events led to the formation of, amongst others, the GNU project and the copyleft movement, which sought to liberate software development from the restraints imposed by copyright. The free software movement prioritised four fundamental freedoms for users: the freedom to run, study, modify, and distribute software. While the term 'free software' was initially used, the label 'open source' was popularized in the late 1990s, emphasising the practical benefits of collaboration and source code sharing. Over time, the movement expanded beyond software, influencing other domains such as hardware, data, and content, and has been integral to the development of countless technologies and platforms that underpin the modern digital age. The initial concerns of the advocates of free software are still relevant today, as the open source definition, while overlapping in many areas, did not always guarantee these initial freedoms, especially as industries found ways to monetize open source without granting users all these rights. On the other hand, ethical open software and AI developers explicitly reject the 'freedom' to run programs for biased or harmful purposes.

With this brief characterisation, we can already begin to see some similarities but also differences between open source movements for software and for AI systems. Historically, the trajectory of enclosure – referring to previously public or shared resources – is at least similar, with much of AI research initially done in an academic setting and subsequently picked up and applied to generate profits. Yet the open source AI movement is complicated by the move of large tech companies investing in AI systems which are made easily accessible to others for their use and integration into applications. These AI systems are also called 'open' for the ease of access, although – and this distinction is important – the models released in such a way usually do not share the source code or allow reconfigurations by others. Open Future has discussed these issues in detail in their take on the Paradox of Open.[17]

Currently the Open Source Initiative is crafting a general definition of open source AI.[18] This definition has not yet become very concrete, yet it centres around licenses that make the study, use, modification and sharing of AI systems and their components by others possible. But the lack of a concrete definition currently has given rise to a variety of different conceptions. Sometimes, what is meant by open source AI includes both the algorithm and the training data; other times, the training

[17] See https://paradox.openfuture.eu/
[18] See https://opensource.org/deepdive/drafts/ and https://opensource.org/deepdive/drafts/the-open-source-ai-definition-draft-v-0-0-3/

data is not considered necessary for an AI system to be open source. Yet another approach is that the trained model is made available, without granting access to the algorithm and/or the data used to train it. What is most often seen as 'open' are AI models which can be used via API access, but these need not be open source in the more expanded view.

A common set of tools

Licensing has been and continues to be a preferred mechanism for ensuring free use of software, including AI systems. Several licenses are available (both restrictive and permissive), but permissive open source licenses are currently the most popular.[19] Open RAIL (Responsible AI Licenses) seek to integrate responsible practices into licensing and have become relatively widespread amongst the restrictive open source licenses.

In general, licenses try to bring the logic used for open source software and code into the framework of AI systems; that is, they offer control over their use through contractual means, effectively securing a legal right to prevent usage under specific conditions. How effective the legal right to prevent usage actually is will be put to the test by the rise of AI systems trained on copyrighted data, which we will discuss next. Another approach for more transparency considered by Wikimedia is governance is governance of AI systems through model cards, allowing the public to engage with how the model is made and suggest edits in the same collaborative vein as with Wikipedia entries.[20] However, a different set of tools also needs to be taken into account, namely the platforms used to share open source AI systems. The rise of HuggingFace as a well-known, central distribution point has made it vastly easier to share AI systems and to draw on the work of others for evaluating and developing applications. It also facilitates research about the kind of AI systems developed in different fields and allows for others to contribute to development. Yet its funding sources indicate that the openness of HuggingFace, as a platform, could be less than permanent and potentially a strategy of enclosure akin to similar marketplace platforms.

What has been sidestepped so far are the infrastructures required for developing open source AI systems, especially those with high computational demands. While it should not be taken for granted that all AI systems require large, indeed ever-increasing, amounts of computing power and data to provide good outputs, the currently popular methods for AI models do rely on computational infrastructure that is in most cases unavailable to research outside of industry.

The question of data

As with all Machine Learning-based AI systems, open source AI systems are heavily reliant on the data used to train them (and all AI systems rely on data for testing). The difficulties start with obtaining datasets which can be used for AI development. Smaller initiatives often have difficulties obtaining sufficiently large datasets for training and testing their systems - this can be due to API restrictions (for example on social media data), to copyright or GDPR considerations, or to difficulties in labelling the data. In tandem with the increasing pushback of copyright holders in connection with generative AI, data available for scraping is quickly dwindling. At the same time, AI-generated content is increasingly spawning all over the Internet, with severe consequences for the

[19] See https://openfuture.pubpub.org/pub/growth-of-responsible-ai-licensing/release/2
[20] https://arxiv.org/abs/1810.03993

usability of online content for the development of generative AI systems: as reported by Carl Franzen,[21] there is evidence that the output of generative AI systems deteriorates if trained on data that consists of too much AI-generated content.[22] Existing copyleft initiatives, for instance around the Creative Commons licenses, could collaborate in the development of open source AI systems – that is, if the computational infrastructure were also available to them or their users.

But open source AI and data intersect in another matter, namely in the issue of licensing. As software, copyright applies to AI systems, hence the focus on licensing that allow open usage (whether restricted or unrestricted). We effectively have analogies of legal frameworks applied to content, which was subsequently extended to software and now becomes applicable to AI systems. The implications of this still need to be explored – for instance, does that mean that open source AI systems should necessarily have open datasets? To which point do open source AI systems need to be open in order to allow for the pursuit of community values?

Points of power

Open source AI systems do exist, but mostly have restrictive licenses, e.g. Apache 2.0. Generally, they are open to varying degrees: most make the model available, without opening up the algorithm or the training data. A famous example is LLaMA, an LLM developed by Meta which was leaked and immediately taken up for experimentation by developers and researchers; yet the non-disclosure of training data (amongst other things) puts the actual openness of the model in questionable. OpenAI's Whisper, a speech-to-text model, and Jukebox, a song generation model, are two of the few models actually released with open licenses from this house. Mistral.ai is one of the European open source developers focusing on LLMs, while EleutherAI's research and releases of open source datasets and LLMs has shifted to increasing interpretability and what they term 'alignment' of AI systems. BLOOM, an open source LLM, was developed by AI researchers and a consortium of companies to allow wider access to LLM capabilities. For many open source developers, affording compute infrastructure and access to data is the biggest challenge. Therefore, funding can be a challenge and is frequently obtained from cloud providers or other big tech industry. In addition to pre-trained models, there are also a large number of open source libraries and other components publicly shared, to make the development of AI systems possible, which are owned by for-profit companies or foundations alike.

But not all open source development faces funding challenges; some companies engage in open source-ish development strategies. Usually, proprietary development ensures the economic viability of products by enclosing them and making use of the system conditional on some exchange of value, be it monetary, expressed in data, or in other forms. However, an equally attractive strategy might be the reconfiguration of others' systems in a way which creates dependence on the developer, thereby exerting an infrastructural pull that can be harnessed (for profit) at opportune moments in time. Effectively, open source with restricted licenses (and the platformisation this strategy effects) provides an opportunity for companies to create revenue streams from standardised services that have become ubiquitous. For instance, the base source code of one of Alphabet's most commonly used services, Android OS, is available under an open

[21] https://venturebeat.com/ai/the-ai-feedback-loop-researchers-warn-of-model-collapse-as-ai-trains-on-ai-generated-content/
[22] https://arxiv.org/pdf/2305.17493v2.pdf

source license, but it is only made public once a new major version has been released. In addition, most Android distributions come with a significant amount of (Google) proprietary software, which (alongside the Android trademark) can only be used through individual licensing agreements. These mechanisms of platformisation point to different pathways towards enclosure for different kinds of market actors and illustrate that openness can also be used to evade regulatory action (in particular related to competition), set de facto standards unilaterally or create other new power constellations, e.g. through financing open source teams. How open source AI will fit into the existing patchwork of strategic enclosure will be interesting to observe.

Value creation

Free and open source movements have historically been closely associated with ideas of commons – that is, communities of practice collaborating to manage shared resources. In a context where AI systems are increasingly embedded into a variety of contexts and products, open source AI systems promise avenues of scrutiny by the public and ways for broader engagement with the development of AI systems, both at a technical and at a political level.

But there are significant value differences between open source movements, ranging from permissionless innovation to actively creating counterpoints to dependencies on large tech companies. One of the questions raised by open source AI systems relates to their wide availability: with generative AI systems being widely available, it becomes easier to automate misinformation as well as many forms of fraud. These are essentially questions of societal control over how technologies are used, and for which purposes; like we have driver's licenses for using automotive technology, but require varied levels of clearance before operating nuclear device, these questions problematise who is fit for using AI technologies.

In the meantime, the label of 'open' has gained a positive connotation, signifying transparency and inclusion; it falls into a similar line of signification as the efforts around 'democratising' AI. With both meanings, it is important to understand who mobilises them when – and there has been a decided shift by large companies towards communicating (marketing) in terms of community values (which may or may not be put into practice). In this context, it is also important that European political leaders have framed open source AI development as a pathway towards open strategic autonomy – or 'digital sovereignty' –, which (amongst other things) is a geopolitical effort to secure independence from providers of strategically important technologies from certain regions. Also perhaps strategically important – but not at all 'open' in terms of known –is the environmental impact of open source AI systems: although in general the environmental impact of AI systems has not yet been fully established, and is significantly more difficult to establish for smaller companies,[23] open source AI systems promise a lower environmental footprint when considering their distribution as the environmental cost of training is distributed amongst a wider user base.

Additionally, the question of value accrual has not been fully settled. The high computational costs of developing AI systems and the challenges with data, as well as the ubiquitous question of remuneration in an open source/open data context, present challenges to the evolving field of open source AI. With proprietary AI systems, the first court cases are popping up to secure remuneration for copyright holders; it remains to be seen whether the stewardship of outputs, often performed by

[23] https://blog.eleuther.ai/fmti-critique/

gig workers at low cost, will be equally counted as contributions to the functioning of these AI systems. What this might mean in a less-resourced open source context remains to be seen.

Recommended reading

For those interested in staying up-to-date on open source AI, Open(ish) Machine Learning News might be a valuable resource. For a very recent review of open source AI, Demos prepared a policy brief as input to the UK's AI Safety Summit. An excellent audio contribution on open source AI comes from the Distributed AI Research Institute's podcast series, Mystery AI Hype Theater 3000. We also liked this ☐HYPERLINK "https://medium.com/@openforgood/open-source-ai-data-sharing-yes-data-colonialism-no-3062a922de03"Medium article if you want to go deeper into the challenges related to data colonialism in an open source AI context. During the event, a guest shared further links on digital sovereignty and building a European digital public space, which we would like to pass on to you.cle if you want to go deeper into the challenges related to data colonialism in an open source AI context. During the event, a guest shared further links on digital sovereignty and building a European digital public space, which we would like to pass on to you.

Profiles of panelists

Zuzanna Warso[24]

Zuzanna works at the intersection of law, ethics and technology. She is the Director of Research at Open Future. With a PhD in law and an M.A. in English Literature, she has received a number of highly celebrated scholarships, including the Marshall Memorial Fellowship from the German Mashall Fund. Since 2019 she has been acting as an independent expert on research ethics to the European Commission.

Jeannette Gorzala[25]

Jeannette is a celebrated member of the Austrian legal community focused on technology, Legal Community Lead and board member of AI Austria, and the vice-president of the European AI Forum. She acts as Associate Professor at the Woxsen University India and has published extensively on the legal implications of AI, blockchain and other emerging technologies.

Alexander Baratsits[26]

Alexander Baratsits has extensively worked as legal counsel, e.g. at the Institute of Science and Technology Austria, and has been a founding board member of Cultural Broadcasting Archive, an association for the promotion of digital communication (cba.media). In addition, he has been extensively involved in the Austrian Creative Commons community, where he has acted as Chapter Lead since 2020. Currently he is active in establishing a European media platform displayeurope.eu.

Katja Mayer

Katja Mayer is Elise Richter Fellow at the University of Vienna, exploring the interface of science, technology, and society, with a particular emphasis on the politics of open science and data infrastructures. Furthermore, she is senior scientist at the Center for Social Innovation ZSI in Vienna and member of the Uni Vienna research platform Governance of Digital Practices.

[24] https://openfuture.eu/author/zuzanna/
[25] Jeannette Gorzala was unable to join as scheduled and has thus provided comments to.
[26] http://www.baratsits.at/cv/