

# “Democratising AI”: Multiple Meanings, Goals, and Methods

Elizabeth Seger<sup>1</sup>  
Centre for the Governance of AI  
Oxford, UK  
[elizabeth.seger@governance.ai](mailto:elizabeth.seger@governance.ai)

Aviv Ovadya  
Harvard Berkman Klein Centre  
Cambridge, MA  
[aviv@aviv.me](mailto:aviv@aviv.me)

Ben Garfinkel  
Centre for the Governance of AI  
Oxford, UK  
[ben.garfinkel@governance.ai](mailto:ben.garfinkel@governance.ai)

Divya Siddarth  
Collective Intelligence Project  
Oxford, UK  
[divya@cip.org](mailto:divya@cip.org)

Allan Dafoe  
Centre for the Governance of AI  
Oxford, UK  
[allan.dafoe@governance.ai](mailto:allan.dafoe@governance.ai)

## ABSTRACT

Numerous parties are calling for “the democratisation of AI”, but the phrase is used to refer to a variety of goals, the pursuit of which sometimes conflict. This paper identifies four kinds of “AI democratisation” that are commonly discussed: (1) the democratisation of AI use, (2) the democratisation of AI development, (3) the democratisation of AI profits, and (4) the democratisation of AI governance. Numerous goals and methods of achieving each form of democratisation are discussed. The main takeaway from this paper is that AI democratisation is a multifarious and sometimes conflicting concept that should not be conflated with improving AI accessibility. If we want to move beyond ambiguous commitments to “democratising AI”, to productive discussions of concrete policies and trade-offs, then we need to recognise the principal role of the democratisation of AI governance in navigating tradeoffs and risks across decisions around use, development, and profits.

## KEYWORDS

AI Democratisation, AI Governance, Model Sharing, AI Benefits, Misuse of AI

## 1 Introduction

Over the last couple years, discussion of “AI democratisation” has surged. AI companies around the world—such as Stability AI [1], Meta [2], Microsoft [3],

and Hugging Face [4]—are talking about their commitment to democratising AI, but it’s not always clear what they mean. The term “AI democratisation” seems to be employed in a variety of ways, causing commentators to speak past one another when discussing the goals, methodologies, risks, and benefits of AI democratisation efforts. This paper aims to provide a foundation for more productive conversations about democratising AI that move beyond ambiguous commitments.

Sections 2 through 5 describe four different notions of AI democratisation that are currently in use by AI labs—democratisation of AI use, democratisation of AI development, democratisation of AI profits, and democratisation of AI governance. We focus on how the term “AI democratisation” is used by AI labs specifically because of the impact they have on the rate of AI advances, and the influence they currently wield over the use, development, profit, and governance of AI.

Each section is divided into two subsections. The first subsection (x.1) discusses various goals the particular form of democratisation is proposed to achieve and notes conflicts with the goals of other forms of democratisation where they arise. The second subsection (x.2) describes various proposed methods for facilitating the form of democratisation.

Although the four concepts of democratisation we discuss often complement each other, it is important to note that they sometimes conflict. For instance, if the public prefers for access to certain kinds of AI systems to be restricted, then the “democratisation of AI governance” may require access restrictions to be put in place—but enacting these restrictions may hinder the “democratisation of AI

---

<sup>1</sup>**Author Contributions:** ES conceived of most of the presented ideas and frameworks, and drafted most of the manuscript with iterative feedback and support from AO, BG, and AD. AO conceptualised and drafted most of Section 5.2. DS conceptualised Section 4 and provided editing support.

development” for which some degree of AI model accessibility is key.

Section 6 then concludes, driving home the main observation of the paper; the first three forms of democratisation (democratisation of use, development, and profits) are about improving accessibility (to AI or AI derived profits) and are not inherently valuable—though the term “democratisation” misleadingly implies otherwise.

## 2 Democratisation of AI Use

When people speak about democratising some technology, they often refer to democratising its use—that is, making it easier for a wide range of people to use the technology. For example, in the early 2010’s the “democratisation of 3D printers” referred to how 3D printers were becoming much more easily acquired, built, and operated by the general public [5].

The same meaning has been applied to the democratisation of AI. Stability AI, for instance, has been a vocal champion of AI democratisation. The company proudly describes its main product, the image generation model Stable Diffusion, as “a text-to-image model that will empower billions of people to create stunning art within seconds” [6]. Microsoft similarly claims to be undertaking an ambitious effort “to democratize Artificial Intelligence (AI), to take it from the ivory towers and make it accessible to all.” A salient part of its plan is “to infuse every application that we interact with, on any device, at any point in time, with intelligence” [3].

### 2.1 Goals

#### 2.2.1 Distributing Benefits of Use

The most commonly articulated goal of democratising AI use, and that communicated above, is to distribute benefits of AI use for many people to enjoy. Benefits include entertainment value (e.g. generating poems with ChatGPT), health and well-being applications, productivity improvement, and other utility functions (writing code, analysing data, creating art). Many of these benefits might, in turn, be translated into financial gains for those who effectively integrate the AI tools into their workstreams. That said, access to these benefits is often financially gated; AI companies that aim to democratise use of their products generally offer their more powerful tools and support services at a price, by purchase or paid subscription.

Ensuring widespread accessibility to some kinds of AI tools may serve also as an important equalising opportunity, providing people with access to services (e.g. medical advice) they would otherwise find difficult to acquire due

to financial, geographical, or transportation restrictions. Similarly, tools like ChatGPT that can be used to draft formal statements and letters may help people to formally vocalise complaints who might otherwise have felt disenfranchised by poor education or limited grasp of the relevant language.

It is important to recognize, however, that for some AI applications the benefits of making the technology available for anyone to use can be relatively minor while the risks are significant. For example, the circle of individuals who would greatly benefit from access to an AI drug discovery tool is relatively small (mainly pharmaceutical researchers), however these tools can be easily repurposed to discover new toxins that might be used as chemical weapons [7]. This is an instance in which unfettered democratisation of AI use—making an AI tool accessible to all—may not always be desirable.

#### 2.2.2 Receiving Feedback for Better and Safer AI

Another reason given for disseminating AI tools widely is so that developers can gather information about how their products are being used (or misused) in a wider variety of contexts than they would have been able to test, let alone imagine, internally [8]. In turn, that feedback informs improvements that enhance model performance and help guard against any new misuse applications that have emerged.

Importantly, where there are concerns about potential misuse, there is an option to cautiously democratise AI use via a staged release of the product [9]. Incrementally larger and more powerful versions of the model are released allowing time between each stage to evaluate how the AI application is being used and to conduct risk benefit analyses of releasing yet a more powerful version of the model. Feedback and staged release may also help provide time and notice for societies to adapt to the new capabilities and harden vulnerable systems, processes, and institutions. Unfortunately, where the risk and consequences of misuse are expected to be severe, responsible AI deployment may require that access restrictions be placed on certain AI capabilities.

Such restrictions limit the democratisation of AI use, but they are not necessarily a blow to AI democratisation more generally. As will be discussed in Section 5, AI democratisation can also refer to the democratisation of AI governance, which is about introducing democratic processes into decision-making about how AI is used, developed, shared and otherwise regulated. Indeed, one might say that the democratisation of AI use—making AI accessible to be used by everyone—is but one possible

outcome of democratising AI governance. It is the outcome if the demos choose that distributing use is desirable.

## 2.2 Methods

Overall, efforts to democratise AI use—to make AI capabilities more widely accessible—may involve reducing the costs of acquiring and running AI tools (again, this may be done via staged release if there are concerns about misuse), providing accessible services to help users integrate AI models into their work streams (e.g. consulting services), and developing intuitive interfaces to facilitate human-AI interaction without the need for extensive training or technical knowhow. In some regions, democratising AI use may also require improvement to more fundamental infrastructure like internet access [10].

## 3 Democratisation of AI Development

When the AI community talks about democratising AI, they rarely limit their focus to the democratisation of AI use. Much of the excitement is about democratising AI development—that is, helping a wider range of people contribute to AI design and development processes.

### 3.1 Goals

#### 3.1.1 Accelerate AI innovation and Progress

A common theme is that tapping into a global community of potential contributors will accelerate innovation and progress in AI research and development. As Microsoft CTO Kevin Scott explains, “we at tech companies, in Silicon Valley startups, can't possibly imagine all of the interesting and valuable things that can be done with [AI], and that we need not just tens or hundreds of thousands of people who work at these companies to be building things but we need hundreds of millions or billions of people able to harness the power of these machines” [11].<sup>2</sup>

It should not be assumed, however, that accelerating AI progress is always desirable. As AI research progresses and AI capabilities improve, we should expect the potential consequence of harms, misuse, or misalignment to also become more severe [12], [13]. The implementation of necessary policy and interventions to ensure safe and responsible AI development going forward may struggle to keep up with unbridled progress, and so there may be a case for exercising some restraint [14].

#### 3.1.2 Cater to Diverse Interests and Needs

Calls for the democratisation of AI development also respond to concern that a small number of leading AI labs

monopolise control over AI development and that those labs employ a narrow demographic of developers. The worry is that the AI products, which are deployed globally, consequently perform disparately for users of different ethnic, geographic, cultural, professional, and financial backgrounds [15]–[18].

Enabling more people to participate in AI design and development processes may help facilitate the development of AI applications that cater to more diverse interests and needs [19]. This is one reason Stability AI offers for its decision to open source Stable Diffusion—meaning that the company allows anyone to download, modify, or build on the Stable Diffusion model on their own computer so long as they agree to the terms of use. CEO Emad Mostaque advocates that “everyone needs to build [AI] technology for themselves.... It's something that we want to enable because nobody knows what's best [for example] for people in Vietnam besides the Vietnamese” [1].<sup>3</sup> The company motto reads “AI by the people, for the people” [20].

But again, it should not be assumed that the diffusion of AI development is universally desirable. Open source sharing in particular may enable more numerous and diverse contributions, but it also opens a door for malicious use and model modification, and controls are difficult to enforce [21].

#### 3.1.3 External Evaluation

Third, many argue that involving more people (e.g. academics, individual developers, smaller labs) in AI development processes provides a critical external evaluation and auditing mechanism. By making models accessible for more people to study, AI labs might distribute auditing duties to a larger and more diverse group of developers than a lab would be able to employ internally. This assumes that more eyes on a model will reveal more flaws leading to safer and more well-aligned technology [8].

### 3.2 Methods

A variety of activities can help enable productive participation in AI design and development processes. Some strategies provide access to AI models and resources to facilitate AI community engagement—e.g. model sharing (3.2.1), providing compute access (3.2.2), project support and coordination (3.2.3). Other strategies help to expand the community of people capable of contributing to AI development processes—e.g. via educational &

---

<sup>2</sup> Quote occurs at 22:30

---

<sup>3</sup> Quote occurs at 13:00

upskilling opportunities (3.2.4) or through the provision of assistive tools (3.2.5).

A key takeaway from this section is that there is much more to AI democratisation—even to the democratisation of AI development specifically—than model dissemination.

### *3.2.1 Model Sharing*

Model sharing involves providing access to AI models including code, model weights and the ability to query, modify, study, or otherwise examine the model.

While model sharing is needed to enable external study and auditing of AI models, it also increases the opportunity for model misuse by malicious actors. In some contexts, for higher risk capabilities, it may therefore be wise to limit model accessibility [22].

That said, model sharing is not an all or nothing activity [23]. Rather, model sharing options range from full open source sharing (all aspects of the system are downloadable for the public) to fully closed (only a select group of developers may even know the model exists) [21], [22], [24]. In the middle there are options for gated access, hosted model access, cloud-based or API access, and downloadable access with some model components withheld. In some of these middle options some degree of advantage from external study and auditing may be maintained while risk of misuse might be reduced.

These options should not be interpreted naively with respect to their stated intent, but with realism about their likely impacts. Google Research published a paper on a technique for style cloning in generative art, and chose to not release any code, citing the potential “societal impact” risk that “malicious parties might try to use such images to mislead viewers” [25]. However, in a mere 11 days one person was able to reproduce the technique to run on Stable Diffusion which they then chose to open source [26]. Similarly, Meta chose to restrict access to the weights of its large language model LLaMa to academic researchers and others on a case by case basis, “to maintain integrity and prevent misuse” [27]. However, a week later the weights (predictably) were leaked and are now available publicly on a torrent [28]. In both these cases, realism is needed in assessing the likely impact of a nominally more restrictive model sharing policy.

### *3.2.2 Improving Compute Access (and other technical infrastructure)*

Large AI models require significant compute power to run. Accordingly, democratising development may also require improvements to compute access. Developers might, for

instance, offer cloud computing services or issue grants for computer cluster access to facilitate smaller and less well-resourced groups in working with more powerful models [21]. Alternatively, developers might explore options for providing smaller model versions that require less compute to run. For example, Emad Mostaque describes Stable Diffusion as “a breakthrough in speed and quality [...] that can be run on consumer GPU’s” [6].

Note, however, that restrictions on compute can also be leveraged to help minimise misuse of powerful AI by limiting the ability of prospective malicious actors to build or modify large models [29], [30]. Therefore, like decisions to open source AI models, decisions to provide significant compute resources should involve adequate risk benefit analysis.

Other tech infrastructure that limits participation in AI development in a similar way to compute access include network accessibility (i.e. access to high bandwidth, low latency internet), access to data storage facilities, access to high quality ethically sourced data, and cyber security infrastructure. These all pose significant barriers to participation in AI development in resource constrained countries, barriers which might be lessened through infrastructure investment and/or remote access [10].

### *3.2.3 Project Support and Coordination*

Democratising AI development is not just about providing resources and assuming that people will come. Effective input elicitation often benefits from dedicated project coordination and support.

For example, the BigScience project was a collaborative effort coordinated by the AI startup Hugging Face—another organisation dedicated to “democratising AI”—and funded by the French government to develop the large language model (LLM) BLOOM [31]. BLOOM was developed over the course of a year by a global coalition of over 1000 volunteer AI developers yielding an LLM functional in 46 languages. Similar collective efforts in other domains may also benefit from funding or other resources to support coordination.

### *3.2.4 Educational and Upskilling Opportunities*

Democratisation of AI development can also be furthered by expanding the community of people capable of making contributions to AI design and development processes.

One option toward this end is for governments and large developer labs to invest in making educational and upskilling opportunities more widely available, especially for demographics traditionally underrepresented in AI

developer communities. Investment in computer science and machine learning educational resources is, for instance, seen as an essential step for establishing AI talent pipelines and narrowing the ‘AI divide’ between the Global North and South [10].

### 3.2.5 Assistive Tools

Another option for expanding the community of prospective contributors is to lower barriers to participation in AI development activities by making it easier for people with minimal programming experience and little familiarity with machine learning to partake.

This might be done through the provision of tools that enable those with less experience and expertise to create and implement their own machine learning applications. For example, Microsoft [32], Google [33], H2O [34] and Amazon [35] have developed “no-code” tools that allow people to build models that are personalised to their own needs without prior coding or machine learning experience. In a similar vein, GitHub Copilot (powered by OpenAI Codex) is a generative AI system that can be used by less experienced developers to help write code [36].<sup>4</sup>

## 4 Democratisation of AI Profits

A third sense of “AI democratisation” refers to democratising AI profits—which is about facilitating the broad and equitable distribution of value accrued to organisations that build and control advanced AI capabilities.

The notion is nicely articulated by Microsoft’s CTO Kevin Scott: “I think we should have objectives around real democratisation of the technology. If the bulk of the value that gets created from AI accrues to a handful of companies in the West Coast of the United States, that is a failure” [11].<sup>5</sup> Though DeepMind does not employ “AI democratisation” terminology, CEO Demis Hassabis expresses a similar sentiment. As reported by TIME, Hassabis believes the wealth generated by advanced AI technologies should be redistributed. “I think we need to make sure that the benefits accrue to as many people as possible—to all of humanity, ideally” [38].

### 4.1 Goals

The goal is rather straightforward: equitably distribute profits generated by AI to ensure wealth and advantages conferred by AI improve human well-being across the

board. A few sub-aims are to avoid widening a socioeconomic divide between AI leading and lagging nations [10], to ease the financial burden of job loss to automation, to smooth economic transition in case of rapid growth of the AI industry, and when AI labs are able to voluntarily participate, to provide mechanisms for labs to powerfully demonstrate their commitment to pursuing advanced AI for the common good [39], [40].

## 4.2 Methods

There are a variety of mechanisms by which AI profits might be more widely distributed or “democratised”. Profits might be redistributed, for instance, via philanthropic giving, though philanthropy can be an inconsistent mechanism of wealth redistribution and, if not well-managed, may worsen inequalities and injustices [41].

Another option is for taxation and profit redistribution to be managed directly by the state [42]. For example, the provision of Universal Basic Income (UBI) has been suggested as a wealth distribution mechanism to help compensate for job loss to automation associated with more advanced AI capabilities [39], [43].

There is concern, however, that taxation methods may be insufficient given the potential of monopolised windfall profits to major AI labs. Accordingly, the proposed “Windfall Clause” offers a third, middle-ground approach [40]. AI firms that voluntarily adopt the Windfall Clause would be bindingly obliged to donate a meaningful portion of their profits when the firm’s profits for the year exceed “a substantial fraction of the world’s total economic output” (e.g. at least 1%). Those donations would then go to a “Distributor” charged with finding and funding effective welfare-maximising projects. Distributors might offer grants to philanthropic organisations, invest directly in infrastructure building projects, or direct funds to state governments for further distribution.

## 5 Democratisation of AI Governance

Finally, some discussions about AI democratisation refer to democratising AI governance. AI governance decisions often involve balancing AI-related risks and benefits to determine if, how, and by whom AI should be used, developed, and shared. The democratisation of AI governance is about distributing influence over these decisions to a wider community of stakeholders and impacted populations. As OpenAI CEO Sam Altman writes, “We want the benefits of, access to, and governance of AGI to be widely and fairly shared” [44].

<sup>4</sup> Though there is concern that users are overconfident and write less secure code when relying on AI assistants like Copilot [37].

<sup>5</sup> Quote occurs at 47:30

## 5.1 Goals

The overarching goal of the democratisation of AI governance is to ensure that decisions around questions such as AI usage, development, and profits reflect the will and preferences of the people being impacted [45]. In this sense, democratisation of AI governance arguably supersedes the previously discussed notions of democratisation. Decisions to democratise use, development, and profits derive their acceptability and desirability from the acceptance and desire of those who will be impacted.

Democratic processes such as referenda, citizen assemblies, and public hearings facilitate the representation of diverse and often conflicting beliefs, opinions, and values into decisions about how people and their actions are governed. Importantly, the desired result is not necessarily agreement among constituents that the best decision was made, but legitimacy—a state of acceptance that the decision-making process was fair and well-considered.

### *5.1.1 Reducing Unilateral Decision-making*

Motivation for democratising AI governance often stems from concern that individual tech companies hold unchecked control over the future of a transformative technology and too much freedom to decide for themselves what constitutes acceptable tradeoffs between risks and benefits of choices around AI use, development, and distribution of profits. It is a worry exacerbated by concern that fierce competition between leading AI labs incentivises reckless decision-making [46].

A single actor in control of a powerful technology or resource can cause significant harm with an ill-considered decision. Consider, for example, the avoidable 2010 Deepwater Horizon oil spill in the Gulf of Mexico. It is one of the greatest environmental disasters in history and largely attributed to a series of cost cutting decisions made by BP including failure to implement proper risk control measures [47]. It is reasonable to assume that in the face of fierce competition and massive financial incentive that AI developers are also liable to make rash decisions about model development and release, the potential negative repercussions of which will only grow as AI capabilities improve [12], [13]. We should therefore not, as Stability AI CEO Emad Mostaque puts it, have “a centralised, unelected entity controlling the most powerful technology in the world” [48]. Introducing democratic processes to enable checks and balances or collective decision-making around AI development, use, and release can potentially guard against ill-considered and potentially detrimental moves.

### *5.1.2 Justice and Fairness*

Another commonly articulated goal is to ensure the benefits and burdens of AI development and deployment are distributed justly and fairly.

It is widely documented that AI systems can replicate or even amplify racial and societal injustices [49], for example, through algorithmic bias in hiring [15], facial recognition [16], loan appraisal [18], and recidivism prediction applications [17]. Some AI misuse cases, such as voice cloning-based phishing, may also have a disproportionate effect on some populations over others [50].

Overall, facilitating the participation or representation of a wide array of stakeholders is seen as a crucial step towards mitigating AI associated injustices [51], [52]. It is to work towards a future for AI in which no communities are disproportionately harmed by development and use activities, and in which no communities are unfairly overlooked as possible beneficiaries of AI capabilities.

### *5.1.3 Navigating Complex Normative Challenges*

The implementation of AI systems in public and private applications raises a variety of normative questions. Some are readily agreed upon such as the high-level assertion that human fatality should be avoided. But there are also many “hard normative questions” to which responses will likely differ depending on culture, context and other value priorities [53]. These challenges include, for example, establishing acceptable risk thresholds, interpreting high-level terms like “transparency” and “dignity”, and determining what values should underpin value-aligned AI [45]. Democratic discourse among diverse stakeholders may help distil areas of agreement or areas in which consensus forming practices are likely to be productive [45]. Inversely, and of equal importance, they might identify cases in which finer-grained details of interpretation and implementation can be determined at a context-specific, local level [53].

## 5.2 Methods

Even though it is already a subcategory of AI democratisation, the democratisation of AI governance is itself a broad and multifaceted concept, some forms of which may be more relevant or useful than others depending on the context [54].

One might speak, for instance, about the introduction of democratic processes to high-level AI policy formation at the national or international governance level or about more fine-grained AI design or deployment decisions made

within individual labs [55]. “Democratic processes” can also refer to a variety of methods for eliciting citizen participation, ensuring substantive representation of stakeholder viewpoints, facilitating well-informed deliberation, holding fair and open election processes, or instituting constitutional protections for individuals and minorities [56].

In what follows, we briefly describe a variety of strategies that have been proposed to underpin democratically legitimate decision making about AI.

### *5.2.1 Harnessing Existing Democratic Structures*

Democratic societies already have many tools and infrastructures in place to facilitate democratically legitimate decision-making about a variety of topics through e.g. legislation and regulation or multilateral standards. Harnessing and modifying effective structures already in place avoids redundancy and reinventing the wheel [55]. It has been proposed, for instance, that with some modification we might make use of procedures laid out by the European Union’s standard-setting organisations (SSOs) to establish context-sensitive standards for safe and responsible AI [53]. We might also structure a new AI governing body after the United States’ FDA (Federal Food and Drug Administration) [57].

While such efforts require minimal new infrastructure, they can, however, get bogged down in existing political quagmires, and are only applicable for decisions that remain within the borders of democratic societies.

### *5.2.2 Multistakeholder Bodies*

Given the global impacts of many AI advances, there has been significant interest in the use of more inclusive processes for input and decision-making around AI. One option is through the formation of multistakeholder bodies to convene diverse, international perspectives for the purpose of navigating complex AI governance challenges. For example, the Partnership on AI (PAI)—a global coalition of academic, civil society, industry, and media organisations—has orchestrated initial non-binding agreements on generative AI across some relevant organisations [58]. There have also been proposals for smaller multistakeholder bodies to form the basis of ethical review boards for high risk AI application and model release decisions [24], [59].

### *5.2.3 Participatory Processes*

Another strategy is to employ modern participatory processes to gather input from diverse populations to guide AI governance decisions. Orchestrating large scale public

participation can be cumbersome and costly, so much promising work focuses on exploring technical solutions such as deliberative tools and digital platforms [60] and generative voting applications [61] to improve the practicability and accessibility of participatory AI governance.

A disadvantage with a participatory approach to governance, however, is that those involved are generally self-selected as stakeholders, community members, or participants, and their outputs may thus only have a weak claim to democratic legitimacy.

### *5.2.4 Representative Deliberation*

A strategy for addressing the challenges of both transnational AI impacts and self-selection is the use of representative deliberation [62], [63], building on heavily researched approaches to deliberative democracy [64]. Representative deliberation involves putting AI governance questions to a representative microcosm of the population of an impacted region, or even the global population (selected by sortition, i.e. stratified sampling) thus granting democratic legitimacy.<sup>6</sup> It is an approach increasingly lauded by both governments [65] and multilateral bodies [66] as a valuable modern approach to democracy generally, and which has found footholds even in authoritarian and global contexts under the less threatening frame of deliberative governance [64], [67].

Companies developing AI systems that want to ‘democratise their governance’ can also delegate such decisions to representative deliberations, and often will have the incentive to do so in the face of competing stakeholder pressures [60], [63]. Meta, for example, has quietly run a set of national and transnational pilots [68] to navigate their ‘complex normative challenges’ and have since scaled up to a near-global deliberative process [69]. Twitter had also planned to pilot such processes before its acquisition [70].

Lighter weight variants of representative deliberations that build on the aforementioned modern participatory practices but in a representative fashion (e.g. using sortition) might also be used to provide a level of democratic legitimacy for less complex AI governance questions [60], [61].

---

<sup>6</sup> Multistakeholder and participatory inputs can integrate with representative deliberations, but the ultimate recommendations are decided on by the representative microcosm, thus granting a stronger claim to democratic legitimacy than multistakeholder or participatory approaches alone.

## 6 Conclusion

This paper has outlined four different notions of “AI democratisation”—the democratisation of use, the democratisation of AI development, the democratisation of AI profits, and the democratisation of AI governance—and discussed numerous goals and methods of achieving each.

For the first three forms of democratisation, “democratisation” is used almost synonymously with “increasing accessibility”. The democratisation of AI use and the democratisation of AI development are about making AI systems accessible for everyone to use or to contribute to their development, and the democratisation of AI profits is about distributing access to profits accrued through AI development and control. The democratisation of AI governance, however, is about balancing these questions of accessibility with other societal needs and values.

Sometimes decisions to democratise AI use, development, and profits will align with societal preferences (ideally determined through democratic processes), and sometimes those preferences will involve restrictions on access. Such is the case, for instance, with legal restrictions on certain medications, treaty restrictions on nuclear weapons, and regulation of labs containing potential hazards. The same may be true of decisions to restrict access to AI models for development or use purposes if risks of open model access are felt to outweigh the benefits.

We should therefore be wary of using the term “democratisation” too loosely or, as is often the case, as a stand-in for “all things good”. The democratisation of AI use, AI development, and even AI derived profits are not inherently good.<sup>7</sup> Their value is derived from alignment with interests and values of those who will be impacted. As such, where the democratically aligned decision would be to limit accessibility, the democratisation of AI governance takes precedence over the others as the source from which the moral and political value of the “democratisation” terminology is derived.

Perhaps the proper response to this paper, then, is to conclude that “AI Democratisation” is a (mostly) unfortunate term. As it is most commonly used within the AI community it refers to facilitating widespread AI use and development which, alone, holds no moral value. Yet invoking the term “democratisation” tells another story. It holds the hidden assumption that the decision to distribute or make accessible is what a democratic governance process would select. In other words, the term AI

---

<sup>7</sup> Improperly managed profit redistribution schemes may worsen inequalities and injustices [41].

democratisation ultimately refers to the democratisation of AI governance. If by “AI democratisation” all a speaker means is “make available to everyone”, then we would suggest less normatively loaded language (something like “broad accessibility”) be used.

## ACKNOWLEDGEMENTS

*We would like to thank Toby Shevlane, Sammy McKinny, Leonie Koessler, Ben Harack, Markus Anderljung, Lennart Heim, Noemi Dreksler, Anton Korinek, Guive Assadi, and Emma Bluemke for their helpful discussion and feedback on various iterations of this draft.*

## REFERENCES

- [1] E. Mostaque, ‘Emad Mostaque (Stability AI): Democratizing AI, Stable Diffusion & Generative Models - Video’, Oct. 23, 2022. Accessed: Mar. 09, 2023. [Online]. Available: <https://exchange.scale.com/public/videos/emad-mostaque-stability-ai-stable-diffusion-open-source>
- [2] ‘Meta AI is sharing OPT-175B, the first 175-billion-parameter language model to be made available to the broader AI research community.’, May 02, 2022. <https://ai.facebook.com/blog/democratizing-access-to-large-scale-language-models-with-opt-175b/> (accessed Mar. 09, 2023).
- [3] ‘Democratizing AI: For Every Person and Organization’, *Microsoft Features*, Sep. 26, 2016. <https://news.microsoft.com/features/democratizing-ai/> (accessed Mar. 09, 2023).
- [4] ‘About us’, *Hugging Face*. <https://huggingface.co/about> (accessed Mar. 09, 2023).
- [5] J. Reitz, ‘3D Printing Today: Democratization of Technology and Disruptive Innovation Converge’, *3DPrint.com | The Voice of 3D Printing / Additive Manufacturing*, Mar. 19, 2018. <https://3dprint.com/207176/democratization-innovation/> (accessed Mar. 09, 2023).
- [6] ‘Stable Diffusion launch announcement’, *Stability AI*. <https://stability.ai/blog/stable-diffusion-announcement> (accessed Mar. 09, 2023).
- [7] F. Urbina, F. Lentzos, C. Invernizzi, and S. Ekins, ‘Dual use of artificial-intelligence-powered drug discovery’, *Nat Mach Intell*, vol. 4, no. 3, Art. no. 3, Mar. 2022, doi: 10.1038/s42256-022-00465-9.
- [8] D. Jeffries, ‘Let’s Speed Up AI’, *Future History*, Feb. 04, 2023. [https://danieljeffries.substack.com/p/lets-speed-up-ai?utm\\_medium=email](https://danieljeffries.substack.com/p/lets-speed-up-ai?utm_medium=email) (accessed Mar. 09, 2023).
- [9] I. Solaiman *et al.*, ‘Release Strategies and the Social Impacts of Language Models’. arXiv, Nov. 12, 2019. Accessed: Feb. 21, 2023. [Online]. Available: <http://arxiv.org/abs/1908.09203>
- [10] D. Yu, H. Rosenfeld, and A. Gupta, ‘The “AI divide” between the Global North and Global



- South', *World Economic Forum*, Jan. 16, 2023. Accessed: Mar. 09, 2023. [Online]. Available: <https://www.weforum.org/agenda/2023/01/davos23-ai-divide-global-north-global-south/>
- [11] K. Scott, 'Democratizing & Accelerating the Future of AI With Kevin Scott - Video', Oct. 07, 2021. Accessed: Mar. 09, 2023. [Online]. Available: <https://exchange.scale.com/public/videos/democratizing-and-accelerating-the-future-of-ai-with-kevin-scott>
- [12] R. Ngo, L. Chan, and S. Mindermann, 'The alignment problem from a deep learning perspective'. arXiv, Feb. 22, 2023. doi: 10.48550/arXiv.2209.00626.
- [13] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜', in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Virtual Event Canada, Mar. 2021, pp. 610–623. doi: 10.1145/3442188.3445922.
- [14] K. Grace, 'Let's think about slowing down AI', *LessWrong*, Dec. 22, 2022. <https://www.lesswrong.com/posts/uFNgRumrDTpBfQGrS/let-s-think-about-slowing-down-ai> (accessed Mar. 09, 2023).
- [15] M. Raub, 'Bots, Bias and Big Data: Artificial Intelligence, Algorithmic Bias and Disparate Impact Liability in Hiring Practices', *Ark. L. Rev.*, vol. 71, p. 529, 2019 2018.
- [16] S. Perkowitz, 'The Bias in the Machine: Facial Recognition Technology and Racial Disparities', *MIT Case Studies in Social and Ethical Responsibilities of Computing*, no. Winter 2021, Feb. 2021, doi: 10.21428/2c646de5.62272586.
- [17] J. A. Mattu Jeff Larson, Lauren Kirchner, Surya, 'Machine Bias', *ProPublica*, May 23, 2016. Accessed: Mar. 09, 2023. [Online]. Available: <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [18] A. Klein, 'Reducing bias in AI-based financial services', Jul. 10, 2020. <https://www.brookings.edu/research/reducing-bias-in-ai-based-financial-services/> (accessed Mar. 09, 2023).
- [19] J. Yang and T. Park, 'Methods for Inclusion: Expanding the Limits of Participatory Design in AI', *Partnership on AI*, Nov. 19, 2020. <https://partnershiponai.org/methodsforinclusion/> (accessed Mar. 10, 2023).
- [20] 'Stability AI', *Stability AI*, Mar. 07, 2023. <https://stability.ai> (accessed Mar. 09, 2023).
- [21] I. Solaiman, 'The Gradient of Generative AI Release: Methods and Considerations'. arXiv, Feb. 05, 2023. Accessed: Feb. 15, 2023. [Online]. Available: <http://arxiv.org/abs/2302.04844>
- [22] A. Ovadya and J. Whittlestone, 'Reducing malicious use of synthetic media research: Considerations and potential release practices for machine learning'. arXiv, Jul. 28, 2019. Accessed: Sep. 30, 2022. [Online]. Available: <http://arxiv.org/abs/1907.11274>
- [23] T. Shevlane, 'Structured access: an emerging paradigm for safe AI deployment'. arXiv, Apr. 11, 2022. Accessed: Feb. 15, 2023. [Online]. Available: <http://arxiv.org/abs/2201.05159>
- [24] P. Liang, R. Bommasani, K. Creel, and R. Reich, 'The Time Is Now to Develop Community Norms for the Release of Foundation Models', *Stanford University Human-Centered Artificial Intelligence; Center for Research on Foundation Models*. <https://crfm.stanford.edu/2022/05/17/community-norms.html> (accessed Mar. 02, 2023).
- [25] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman, 'DreamBooth: Fine Tuning Text-to-Image Diffusion Models for Subject-Driven Generation'. arXiv, Aug. 25, 2022. Accessed: Mar. 15, 2023. [Online]. Available: <http://arxiv.org/abs/2208.12242>
- [26] Xavier, 'Dreambooth on Stable Diffusion'. Mar. 15, 2023. Accessed: Mar. 15, 2023. [Online]. Available: <https://github.com/XavierXiao/Dreambooth-Stable-Diffusion>
- [27] 'Introducing LLaMA: A foundational, 65-billion-parameter language model', *Meta AI*, Feb. 24, 2023. <https://ai.facebook.com/blog/large-language-model-llama-meta-ai/> (accessed Mar. 15, 2023).
- [28] J. Vincent, 'Meta's powerful AI language model has leaked online — what happens now?', *The Verge*, Mar. 08, 2023. <https://www.theverge.com/2023/3/8/23629362/meta-ai-language-model-llama-leak-online-misuse> (accessed Mar. 15, 2023).
- [29] L. Heim and M. Anderljung, 'GovAI Response to the Future of Compute Review - Call for Evidence'. Accessed: Mar. 09, 2023. [Online]. Available: <https://www.governance.ai/research-paper/future-of-compute-review-call-for-evidence>
- [30] L. Heim and M. Anderljung, 'Submission to the Request for Information (RFI) on Implementing Initial Findings and Recommendations of the NAIRR Task Force | GovAI'. Accessed: Mar. 09, 2023. [Online]. Available: <https://www.governance.ai/research-paper/submission-nairr-task-force>
- [31] M. Heikkila, 'BLOOM: Inside the radical new project to democratize AI | MIT Technology Review', *MIT Tech Review*, Jul. 12, 2022. Accessed: Mar. 09, 2023. [Online]. Available: <https://www.technologyreview.com/2022/07/12/1055817/inside-a-radical-new-project-to-democratize-ai/>
- [32] 'What is automated ML? AutoML - Azure Machine Learning', *Microsoft Learn*, Feb. 24, 2023. <https://learn.microsoft.com/en-us/azure/machine-learning/concept-automated-ml> (accessed Mar. 09, 2023).
- [33] 'Google Cloud AutoML - Train models without ML expertise', *Google Cloud*. <https://cloud.google.com/automl> (accessed Mar. 09, 2023).
- [34] 'H2O Driverless AI'. <https://h2o.ai/platform/ai-cloud/make/h2o-driverless-ai/> (accessed Mar. 09, 2023).

- [35] ‘No-code machine learning - Amazon Web Services’, *Amazon Web Services, Inc.* <https://aws.amazon.com/sagemaker/canvas/> (accessed Mar. 09, 2023).
- [36] T. Warren, ‘GitHub’s AI-powered Copilot will help you write code for \$10 a month - The Verge’, *The Verge*, Jun. 21, 2022. Accessed: Mar. 09, 2023. [Online]. Available: <https://www.theverge.com/2022/6/21/23176574/github-copilot-launch-pricing-release-date>
- [37] N. Perry, M. Srivastava, D. Kumar, and D. Boneh, ‘Do Users Write More Insecure Code with AI Assistants?’ arXiv, Dec. 16, 2022. doi: 10.48550/arXiv.2211.03622.
- [38] B. Perrigo, ‘DeepMind CEO Demis Hassabis Urges Caution on AI | TIME’, *TIME*, Jan. 12, 2023. Accessed: Mar. 09, 2023. [Online]. Available: <https://time.com/6246119/demis-hassabis-deepmind-interview/>
- [39] K. Miller, ‘Radical Proposal: Universal Basic Income to Offset Job Losses Due to Automation’, *Stanford HAI*, Oct. 20, 2021. <https://hai.stanford.edu/news/radical-proposal-universal-basic-income-offset-job-losses-due-automation> (accessed Mar. 09, 2023).
- [40] C. O’Keefe, P. Cihon, B. Garfinkel, C. Flynn, J. Leung, and A. Dafoe, ‘The Windfall Clause: Distributing the Benefits of AI’, Future of Humanity Institute, University of Oxford, 2020. Accessed: Mar. 10, 2023. [Online]. Available: <https://www.fhi.ox.ac.uk/wp-content/uploads/Windfall-Clause-Report.pdf>
- [41] R. Reich, *Just giving: why philanthropy is failing democracy and how it can do better*. Princeton, New Jersey: Princeton University Press, 2018.
- [42] S. Altman, ‘Moore’s Law for Everything’, *Moore’s Law for Everything*, Mar. 16, 2021. <https://moores.samaltman.com/> (accessed Mar. 15, 2023).
- [43] ‘Barack Obama Talks AI, Robo Cars, and the Future of the World: The president in conversation with MIT’s Joi Ito and WIRED editor-in-chief Scott Dadich.’, *Wired*. Accessed: Mar. 15, 2023. [Online]. Available: <https://www.wired.com/2016/10/president-obama-mit-joi-ito-interview/>
- [44] S. Altman, ‘Planning for AGI and beyond’, *OpenAI*. <https://openai.com/blog/planning-for-agi-and-beyond> (accessed Mar. 14, 2023).
- [45] I. Gabriel, ‘Artificial Intelligence, Values, and Alignment’, *Minds & Machines*, vol. 30, no. 3, pp. 411–437, Sep. 2020, doi: 10.1007/s11023-020-09539-2.
- [46] K. Piper, ‘Are we racing toward AI catastrophe?’, *Vox*, Feb. 09, 2023. <https://www.vox.com/future-perfect/23591534/chatgpt-artificial-intelligence-google-baidu-microsoft-openai> (accessed Mar. 09, 2023).
- [47] S. Goldenberg, ‘BP cost-cutting blamed for “avoidable” Deepwater Horizon oil spill’, *The Guardian*, Jan. 06, 2011. Accessed: Mar. 15, 2023. [Online]. Available: <https://www.theguardian.com/environment/2011/jan/06/bp-oil-spill-deepwater-horizon>
- [48] K. Roose, ‘A Coming-Out Party for Generative A.I., Silicon Valley’s New Craze’, *The New York Times*, Oct. 21, 2022. Accessed: Mar. 09, 2023. [Online]. Available: <https://www.nytimes.com/2022/10/21/technology/generative-ai.html>
- [49] S. U. Noble, *Algorithms of oppression: how search engines reinforce racism*. New York: New York University Press, 2018.
- [50] P. Verma, ‘They thought loved ones were calling for help. It was an AI scam.’, *Washington Post*, Mar. 06, 2023. Accessed: Mar. 09, 2023. [Online]. Available: <https://www.washingtonpost.com/technology/2023/03/05/ai-voice-scam/>
- [51] A. Zimmermann, E. Di Rosa, and H. Kim, ‘Technology Can’t Fix Algorithmic Injustice’, *Boston Review*, Jan. 09, 2020. Accessed: Mar. 09, 2023. [Online]. Available: <https://www.bostonreview.net/articles/annette-zimmermann-algorithmic-political/>
- [52] P. Kalluri, ‘Don’t ask if artificial intelligence is good or fair, ask how it shifts power’, *Nature*, vol. 583, no. 7815, pp. 169–169, Jul. 2020, doi: 10.1038/d41586-020-02003-2.
- [53] J. Laux, S. Wachter, and B. Mittelstadt, ‘Three Pathways for Standardisation and Ethical Disclosure by Default under the European Union Artificial Intelligence Act’. Rochester, NY, Feb. 20, 2023. doi: 10.2139/ssrn.4365079.
- [54] H. S. Sætra, H. Borgebund, and M. Coeckelbergh, ‘Avoid diluting democracy by algorithms’, *Nat Mach Intell*, vol. 4, no. 10, pp. 804–806, Sep. 2022, doi: 10.1038/s42256-022-00537-w.
- [55] J. Himmelreich, ‘Against “Democratizing AI”’, *AI & Soc*, Jan. 2022, doi: 10.1007/s00146-021-01357-z.
- [56] P. Chiochetti, ‘Democratic legitimacy’, *Reveu de l’euro*, 2017, doi: 10.25517/RESUME-7XN4KF9-2017.
- [57] A. Tutt, ‘An FDA for Algorithms’. Rochester, NY, Mar. 15, 2016. doi: 10.2139/ssrn.2747994.
- [58] ‘PAI’s Responsible Practices for Synthetic Media’, Partnership on AI, Feb. 2023. Accessed: Mar. 10, 2023. [Online]. Available: [https://partnershiponai.org/wp-content/uploads/2023/02/PAI\\_synthetic\\_media\\_framework.pdf](https://partnershiponai.org/wp-content/uploads/2023/02/PAI_synthetic_media_framework.pdf)
- [59] J. Schuett, A. Reuel, and A. Carlier, ‘AI ethics boards: Design considerations for reducing risks from AI’, Forthcoming.
- [60] ‘Introducing the Collective Intelligence Project: Solving the Transformative Technology Trilemma through Governance R&D’, *The Collective Intelligence Project*, 2023. <https://cip.org/whitepaper> (accessed Mar. 10, 2023).
- [61] A. Ovadya, ‘“Generative AI” through Collective Response Systems’. arXiv, Feb. 01, 2023. Accessed: Feb. 09, 2023. [Online]. Available: <http://arxiv.org/abs/2302.00672>
- [62] F. Carugati, ‘A Council of Citizens Should Regulate Algorithms’, *Wired*. Accessed: Mar. 10, 2023.

- [Online]. Available:  
<https://www.wired.com/story/opinion-a-council-of-citizens-should-regulate-algorithms/>
- [63] A. Ovadya, ‘Towards Platform Democracy: Policymaking Beyond Corporate CEOs and Partisan Pressure’, Belfer Center for Science and International Affairs, Oct. 2021. Accessed: Mar. 10, 2023. [Online]. Available:  
<https://www.belfercenter.org/publication/towards-platform-democracy-policymaking-beyond-corporate-ceos-and-partisan-pressure>
- [64] J. S. Dryzek *et al.*, ‘The crisis of democracy and the science of deliberation’, *Science*, vol. 363, no. 6432, pp. 1144–1146, Mar. 2019, doi: 10.1126/science.aaw2694.
- [65] OECD, ‘Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave’, OECD, Jun. 2020. doi: 10.1787/339306da-en.
- [66] ‘European Citizens’ Panels - Conference on the Future of Europe’.  
<https://futureu.europa.eu/en/assemblies/citizens-panels> (accessed Mar. 13, 2023).
- [67] ‘Report of the 2021 Global Assembly on the Climate and Ecological Crisis: Giving everyone a seat at the global governance table’, Global Assembly, Nov. 2022. Accessed: Mar. 13, 2023. [Online]. Available:  
<https://globalassembly.org/resources/downloads/GlobalAssembly2021-FullReport.pdf>
- [68] ‘Deliberative democracy in action: A closer look at our recent pilot with Meta’, *The Behavioural Insights Team*.  
<https://www.bi.team/blogs/deliberative-democracy-in-action/> (accessed Mar. 13, 2023).
- [69] B. Harris, ‘Improving People’s Experiences Through Community Forums’, *Meta*, Nov. 16, 2022. <https://about.fb.com/news/2022/11/improving-peoples-experiences-through-community-forums/> (accessed Mar. 13, 2023).
- [70] A. Ovadya, “‘Platform Democracy’—a very different way to govern big tech”, *Reimagining Technology*, Nov. 16, 2022. <https://aviv.substack.com/p/platform-democracy-a-different-way-to-govern> (accessed Mar. 13, 2023).