

---

# REDUCING MALICIOUS USE OF SYNTHETIC MEDIA RESEARCH: CONSIDERATIONS AND POTENTIAL RELEASE PRACTICES FOR MACHINE LEARNING

---

**Aviv Ovadya**

The Thoughtful Technology Project  
aviv.ovadya@thoughtfultech.org

**Jess Whittlestone**

Leverhulme Centre for the Future of Intelligence  
University of Cambridge  
jlw84@cam.ac.uk

## ABSTRACT

The aim of this paper is to facilitate nuanced discussion around research norms and practices to mitigate the harmful impacts of advances in machine learning (ML). We focus particularly on the use of ML to create “synthetic media” (e.g. to generate or manipulate audio, video, images, and text), and the question of what publication and release processes around such research might look like, though many of the considerations discussed will apply to ML research more broadly. We are not arguing for any specific approach on when or how research should be distributed, but instead try to lay out some useful tools, analogies, and options for thinking about these issues.

We begin with some background on the idea that ML research might be misused in harmful ways, and why advances in synthetic media, in particular, are raising concerns. We then outline in more detail some of the different paths to harm from ML research, before reviewing research risk mitigation strategies in other fields and identifying components that seem most worth emulating in the ML and synthetic media research communities. Next, we outline some important dimensions of disagreement on these issues which risk polarizing conversations.

Finally, we conclude with recommendations, suggesting that the machine learning community might benefit from: working with subject matter experts to increase understanding of the risk landscape and possible mitigation strategies; building a community and norms around understanding the impacts of ML research, e.g. through regular workshops at major conferences; and establishing institutions and systems to support release practices that would otherwise be onerous and error-prone.

**Keywords** synthetic media · research practices · release practices · societal impacts · negative impacts · machine learning · malicious use · dual-use · deepfakes · fake news

## 1 Introduction

Technological advances can result in harm to both individuals and societal structures, through accidents, unintended consequences, and malicious use – even as the same advances provide incredible benefits. Concern about harms resulting from advances in machine learning (ML) have risen dramatically in recent years, with a growing community of researchers and practitioners doing crucial work to address issues such as the fairness, accountability, and transparency of deployed systems [1]. In this paper, we focus primarily on the ways that advances in ML research might be deliberately misused by malicious actors (which we refer to as “mal-use”) [2], though we also touch on other unintended consequences.<sup>1</sup>

---

<sup>1</sup>Unintended consequences are also particularly salient for the case of synthetic video and audio, where the simple existence of the technology, regardless of its use, can allow bad actors to claim that evidence of e.g. corruption or war crimes were synthesized in order to avoid accountability. For example, allegations of videos being faked have been used to justify a coup in Gabon [3], and exculpate a cabinet minister in Malaysia [4].

One specific example of where mal-use concerns have been raised recently is around the use of ML in *synthetic media*<sup>2</sup>: the manipulation and generation of increasingly realistic audio, video, images, and text. It is now possible to produce synthetic images of faces that are almost indistinguishable from photographs [5]. It is becoming easier to manipulate existing videos of people: replacing a person’s facial expressions or movements with those of another person [6, 7], and/or overlaying synthesized speech which imitates a person’s voice [8]. Language models have advanced to the point where, given a sentence prompt, they can write articles that appear at least superficially convincing [9]. These advances could be used to impersonate people, sway public opinion, or more generally spread doubt about the veracity of all media. Modern synthetic media is in fact already being used for harm: face-swapping tools are being used to harass journalists [10], synthetic voices are being used for financial crimes [11], and synthetic faces have allegedly been used for espionage [12].

These examples bring into focus the need for humility around what we don’t know concerning potential impacts of technology, and perhaps suggests that we must develop better systems for understanding such impacts. Related concerns have sparked debate about responsible research practices in ML [13], and in particular whether sometimes it is appropriate to withhold open publication of some aspects of this research.

We begin with some background on the idea that ML research might be misused in harmful ways, and why advances in synthetic media, in particular, are raising concerns. We then review research risk mitigation strategies in other fields and identifying components that may be worth emulating in the ML and synthetic media research communities. Finally, we outline some important dimensions of disagreement on these issues which risk polarizing conversations, before concluding with some recommendations for research and practice.

## 2 How research can lead to harm

How might advances in ML and synthetic media end up resulting in harm? We begin by spelling this out in more detail: different ways that research might empower malicious actors, and some possible paths to harm from this.

### 2.1 Types of hazard

We can distinguish several different types of “information hazard” that might result from machine learning research in general.<sup>3</sup>

- **Product hazard:** *Research produces software that can be directly used for harm* (e.g. rootkits for computer hacking). Product hazards increase the likelihood of mal-use by adversaries with minimal technical capabilities (e.g. ‘script kiddies’ who use existing programs or scripts to hack into computers but lack the ability to write their own, or less sophisticated information warfare operations), or those who have weak motivations for harm (e.g. online hacking or deception just ‘for fun’).
- **Data hazard:** *research produces detailed information or outputs which if disseminated, create risk of use for harm* (e.g. easy nuclear blueprints; models, training data or code for harmful software). Data hazards increase the likelihood of mal-use by adversaries with some technical capabilities, but without e.g. access to high-quality researchers.
- **Attention hazard:** *research which directs attention towards an idea or data that increases risk* (e.g. the idea that it is possible to use voice-cloning for phishing). Attention hazards increase the likelihood of mal-use by adversaries who may not have realized that their objectives can be aided by new technologies.

An important thing to note is that potential mitigations will be *different for each type of hazard – and potentially even in conflict*. One way to mitigate attention hazards is to be very careful about talking to media organizations and others with large reach about ways that research advances could be used maliciously (such as in voice cloning, for example). At the same time, raising wider concern about malicious use cases of ML progress might be exactly what is needed to incentivize mitigations for data hazards or product hazards (e.g. some public concern may be required to ensure that tech platforms prioritize developing technology to identify voice cloning, or for telecoms to build mitigating infrastructure to make phone number spoofing harder).

<sup>2</sup>Not all synthetic media techniques involve ML. For example, traditional computer graphics techniques are also used for image and video manipulation. For simplicity of language, we will not consider these distinctly as the considerations for research practices are very similar.

<sup>3</sup>This is based loosely on a taxonomy from [14]

## 2.2 The path to harm

But how might these different types of ‘hazard’ actually lead to real-world harms? It is worth connecting the dots here between the theoretical potential for mal-use, and what actually makes significant harm more likely.

Below we talk through some factors influencing whether a capability leads to sustained mal-use in practice. We use artificial voice cloning as an illustrative example, as a relatively new capability with many useful applications (e.g. in voice translation and audio editing) but also significant potential for mal-use (e.g. in scams, political propaganda, and market manipulation).

### 1. *Awareness: Do actors with malicious intent know about a capability and believe it can help them?*

We can break this down into:

- **Attention of adversaries:** Are malicious actors likely to realize that they could use a new capability to further their ends? If adversary groups are already using closely related methods, this is much more likely: for example, if edited voice clips are already being used for political manipulation, groups doing this are more likely to pay attention to demonstrations of voice cloning.
- **‘Convincibility’ of those with resources:** Are there compelling arguments, perhaps by authoritative third parties, for the effectiveness of new capabilities? For example, a scammer who realizes that voice cloning is useful might need to be able to convince a superior that this technology is effective enough to justify the costs and overcome institutional inertia.

### 2. *Deployment: How difficult is it for adversaries to weaponize this capability in practice?*

For a capability to be deployed for malicious purposes, adversaries not only need to be aware but to have the necessary skills and resources to productize and weaponize the capability. This isn’t a binary – e.g. having ML expertise vs. not – but rather many different factors will influence how easy a capability is to weaponize. At the extreme, we might have a product which can be immediately used by anyone, regardless of technical capability (such as free to use voice cloning software).

Factors that influence the ease of deployment for mal-use include:

- **Talent pipelines:** How difficult is it to source someone who can apply a new capability for the desired use case? (e.g. do malicious actors need someone with machine learning experience, programming experience, or can they just use a program directly to achieve their goals?) [12].
- **Reproducibility:** How difficult is it to reproduce a capability given the information available? (e.g. is it easy to replicate a voice cloning capability given the available papers, models, code, etc.?)
- **Modifiability:** How difficult is it to modify or use a system in order to enable mal-use? (e.g. if a voice cloning product makes it difficult to clone a voice without consent or watermarks, how hard is it to overcome those limitations?)
- **Slottability:** Can new capabilities be slotted into existing organizational processes or technical systems? (e.g. are there already established processes for phone scams into which new voice generation capabilities can be slotted easily, without any need to change goals or strategy?)
- **Environmental factors:** How does the existing ‘environment’ or ‘infrastructure’ impact the usefulness of the new capability for malicious actors? (E.g. currently, in the US it is easy to ‘spoof’ phone numbers to make it appear like a call is coming from a family member, which could impact the likelihood of voice cloning being weaponized for phone scams.)

Websites now enabling anyone to instantly generate seemingly photorealistic faces are a concrete example of deployment barriers falling away and making mal-use easier. It had been possible for well over a year to generate synthetic images of faces with fairly high quality, but such websites have enabled anyone to do so with no technical expertise. This capability can also immediately slot into existing processes, such as fake account creation. Previously, malicious actors would often use existing photos of real people, which could be identified with reverse image search [15], unlike wholly generated synthetic images.

### 3. *Sustained use: How likely is it that a capability will lead to sustained use with substantial negative impacts?*

Even if adversaries are aware of and able to weaponize some new capability, whether or not this leads to sustained use depends on:

- **Actual ROI:** If malicious actors believe that the return ROI on investment (ROI) for using a capability is low they might not continue to use it in practice. For example, if a form of mal-use is easy to detect, then adversaries might decide it’s not worth the risk or might be shut down very quickly.

- **Assessment of ROI:** If malicious actors have no way of assessing whether new capabilities are helping them better achieve their goals, or if their assessments are flawed, they might not continue to put resources into using those capabilities.

### 2.3 Access ratchets

We can think of this as a kind of progression, from a theoretical capability to scaled-up use in practice. Once a technology has progressed down this path and has become easy to use, and proven to have high ROI for mal-use, it can be much more difficult to address than at earlier stages – we call this the **access ratchet** (like a ratchet, increased access to technology cannot generally be undone). For any capability with potential for mal-use, it is therefore worth thinking about where it currently sits on this progression: how much attention and interest it is receiving; whether it has been weaponized and/or how costly it would be to do so; and whether it’s likely to be, or already in sustained use. This can help us think more clearly about where the greatest risks of mal-use are, and different kinds of interventions that might be appropriate or necessary in a given situation.

Researchers may argue that a capability is unlikely to cause harm since it has not been used maliciously yet. What this doesn’t address is the fact that a capability which has not yet been used maliciously might sit anywhere along this progression, which makes a huge difference to how likely it is to cause harm. For example, Face2Face, a technique for real-time facial reenactment (i.e. changing a person’s expressions in a video), has existed for 3 years but not been developed into any products that can easily be used. This lack of productization makes harmful use vastly less likely, especially given the competition for AI and engineering talent today. It is also worth considering how costly it would be to make a given capability easier to misuse: even the DeepFake application, which is more accessible to non-technical users, is currently resource-intensive to weaponize in practice.

### 2.4 Indirect harms

Sometimes the path to harm from synthetic media research will be fairly direct and immediate: such as a person losing their money, returning to our example of voice cloning being used in financial scams.

But in other cases, improved synthetic media capabilities might cause harm in more complex and indirect ways. Consider the case where misinformation purveyors get hold of sophisticated synthetic media capabilities and use them to win substantial democratic power, which they then use to control narratives further and undermine any mitigation efforts (not an uncommon path from democracy to authoritarianism). We can think about this as a **disinformation ratchet**: the ability to use disinformation to enhance one’s ability to distribute further disinformation; and the opportunity for this type of ratchet can be influenced by new technology impacting media distribution channels and capabilities.

These less direct kinds of harms may be harder to anticipate or imagine, but in the long-run may be much more important – particularly if they influence the future development of technology in ways that undermine our ability to deal with future threats. We suggest that it’s particularly important to consider these kinds of “sociotechnical-path dependencies” as well as more direct and immediate threats, and what kinds of risk mitigation strategies might best address them.

## 3 Mitigating harm through release practices

There is unlikely to be any ‘one size fits all’ solution to mitigating mal-use of ML research: the path to harm will look very different across contexts, and potential harms need to be weighed against benefits which will also vary depending on the area. We therefore need discussion about different approaches to mitigating mal-use: including around what research is conducted in the first place; standards and procedures for risk assessment; and processes for deciding when and how to release different types of research outputs. Here we focus particularly on the latter – how careful release practices might help mitigate mal-use within ML research.

However, this is not to suggest we think release practices are the main or even necessarily the most important component of mitigating mal-use. Another crucial piece is how research directions are chosen and prioritized in the first place. This is challenging because much of ML research often involves developing general capabilities which can then be applied to a variety of different purposes – we can’t simply decide to build only ‘beneficial’ ML capabilities. That aside, we still may be able to say some very general things about the kinds of capabilities that are more likely to be broadly beneficial, or the kinds of problems that should ideally be driving ML research. It is also important to think about what types of research are encouraged/discouraged by conferences, journals, funders, job interviewers and so on.

### 3.1 Challenges to mitigating harm

First, it's worth considering some of the serious challenges to attempting to decrease harm by limiting access to research:

- The **composition problem**: Two independent pieces of research that seem innocent can be combined in ways that enable significant malicious use.<sup>4</sup>
- The **slow drip problem**: Research advancement can be a slow and continuous evolution, where it's difficult to draw the line between research that is dangerous and that which is not.
- The **conflation problem**: Many of the underlying goals of various fields of research (natural language processing, computation photography, etc.) may be directly weaponizable if achieved. For example, the ability to create convincing dialogue can be used to both support people or manipulate people at scale.
- The **defector problem**: Even if researchers in some regions or organizations cooperatively decide not to pursue or publish a particular area of research, those agreements might not be followed by “defectors” who then gain a competitive edge.

These challenges may seem daunting even for those who would advocate for limiting release of some forms of research. They also motivate the development of a nuanced menu of options for release practices, and careful evaluation of the efficacy of whatever measures are chosen. Even without overcoming these challenges, it is possible that release practices could substantially mitigate harm if they impact the rate of deployment of mal-use technology.<sup>5</sup>

### 3.2 A brief tour of analogs in other fields

There is precedent in several other fields of research – including biotechnology and information security – for establishing processes for reducing the negative risks of research and release. A good first step would, therefore, be to look at what we can learn from these fields for the case of ML research. Here we present some promising practices identified from other fields.

A caveat: just because research norms and processes exist in other fields, it does not necessarily mean that they are widely and coherently used in those fields, or that they provide a net positive impact. Evaluating which research practices have been adopted and work well across different fields is out of scope for this short paper, but would certainly be valuable to look into further.

#### 3.2.1 Biosafety

Biosafety processes and principles exist to ensure safe handling of infective microorganisms in biology/biotechnology research [16]. Some key components of biosafety practices include:

- *Procedures*: Steps and rules that must be followed, e.g. for decontamination (including basics such as wearing gloves and shoes).
- *Lab safety officer*: An internal role responsible for enforcing safety.
- *Training*: Learning safety processes via peers/programs.
- *Architecture*: Incorporating safety considerations into building and tool design (e.g. the design of doors and airflow).
- *Audits*: Providing external accountability, usually at random times via local government.
- *Safety level designations*: Different microorganisms classified by risk group (e.g. Ebola is level 4) with different safety procedures for different levels (e.g. level 1 is open bench work, level 4 requires special clothing, airlock entry, special waste disposal, etc.).
- *Safety level definers*: Organisations who determine safety levels, e.g. the Centers for Disease Control and Prevention (CDC).

<sup>4</sup>This might be particularly challenging given the success of transfer learning.

<sup>5</sup>In terms of the ratchet terminology used earlier, delaying release of research could slow down the speed of an ‘access ratchet’ (i.e. slowing down widespread access to a technology), potentially providing enough extra time strengthen a ‘sociotechnical immune system’ that could halt a disinformation ratchet.

### 3.2.2 Computer/Information security

Various practices exist in the field of information security to prevent exploitation of vulnerabilities in important systems. Key components include:

- *OPSEC ('operations security')*: Procedures for identifying and protecting critical information that could be used by adversaries. Includes identification of critical information, analysis of threats, vulnerabilities, and risks, and application of appropriate measures.
- *Architecture*: Use systems that are “secure by design” and so keep you secure automatically where possible.
- *Coordinated/responsible disclosure*: Processes to ensure that exploits which could affect important systems are not publicly disclosed until there has been an opportunity to fix the vulnerability.
- *ISACs/CERTs* (Information Sharing & Analysis Centers/Computer Emergency Response Teams): Disclosure coordination entities.

### 3.2.3 Institutional Review Boards (IRBs)

IRBs are designed to protect human subjects in biomedical and behavioral research (including e.g. clinical trials of new drugs or devices and psychology studies of behavior, opinions or attitudes) [17]:

- *Case dependent scrutiny*: Research proposals are assessed on a case-by-case basis using external expert evaluation, and are determined to either be: (a) exempt (when risks are minimal), (b) expedited (slightly more than minimal risk), or (c) full review (all other proposals).
- *Approval rubrics*: Criteria for approval of research proposals include: having sound research principles to minimize risk to subjects; establishing that risks to subjects are reasonable relative to anticipated benefits; selecting subjects in equitable ways, and avoiding undue emphasis on a vulnerable population.
- *External expert & community evaluation*: Studies are reviewed by people who have expertise in the research and in the impacts of the work (such as community members).
- *Continuous evaluation*: Process can be ongoing, not one-time, with periodic updates: the IRB can suspend or terminate previously approved research.

This is not meant to be exhaustive but demonstrates a variety of systems that have been used to mitigate negative risks of research and release. Other analogs worth exploring include those around nuclear technology, spam detection, classified information, and environmental impact.

### 3.3 Potential release practices

What should ML and synthetic media research emulate from these other fields? Many aspects of these practices and processes may be applicable in ML research, including particularly: external expert evaluation of risks and appropriate responses, coordinated/responsible disclosure, training in responsible research processes, disclosure coordination entities, safety level designations, safety level defining entities, and case-dependent response (depending on safety levels).

Reframing and renaming all of these practices and processes to focus on the release of potentially hazardous ML systems leaves us with the following components that may be needed in ML:

- **Release options**: Different options for release.
- **Release rubric**: Guidelines for when to use each type (decided by case-dependent evaluation).
- **Release rubric processes**: How to do case-dependent evaluation.
- **Release coordination**: Who decides/gets access.
- **Release training**: How to learn processes/norms.
- **Release process entities**: Who manages all of this?

Each of these components can be broken down further; we explore “release options” here as an example.

### 3.4 Release options

The question of how to release research with potential for mal-use is not a binary one: there are many different choices to make beyond simply ‘release’ or ‘don’t release’. Focusing on this binary choice can lead the debate around openness of ML research to become very polarized.

Some important dimensions we might consider when thinking about release strategies include:

- **Content: *What is released***

Potential options include:

- A fully runnable system (with varying power)
- A modifiable system (with varying modifiability)
- Source code (varying versions)
- Training data (varying sizes)
- Trained models (varying strength/fine-tunability/data-needs)
- Paper/concept (varying detail level)
- Harmful use case ideas (varying detail level)

- **Timing: *When it is released***

Potential options include:

- Immediate release
- Timed release: Set a specific time to release components, allowing time for mitigation of any potential harms. This is common in information security.
- Periodic evaluation: Don’t release immediately, but set a time period/intervals (e.g. every 2 months), at which point an evaluation is done to reassess the risk of release given mitigation progress.
- Evented release: Wait to release until some particular type of external event (e.g. someone else replicating or publicizing the same technology).
- Staged release: Release systems of successively increasing levels of power on a fixed timeline, or triggered by external events.

- **Distribution: *Where/Who it is released to***

Potential options include:

- Public access (with varying degrees of publicity)
- Ask for access: Anyone who wants access to data or a system asks and is approved on a case-by-case basis, potentially with specific requirements around use.
- Release safety levels: People and possibly organizations can request to be recognized as ‘safe’, after auditing and approval they gain the ability to access all material at a given safety level.
- Access communities: Research groups developing their own trusted communities through informal processes which all have access to shared repositories.

Within the domain of synthetic media, it’s worth diving deeper into potential mitigations specific to products, models, and demos relevant to that space.

There are a number of mechanisms researchers and companies can use to reduce malicious use from general synthetic media systems that allow e.g. virtual impersonation:

- **Consent:** Requiring consent by those being impersonated.
- **Detectability:** Intentionally not trying to thwart detection.
- **Watermarking:** Embedding context about modifications/original.
- **Referenceability:** Centrally storing all modifications for reference.

It’s important to note that none of these are perfect – they are part of a “defense in depth”. It is also possible to add constraints on synthesis (e.g. ensuring that only particular faces can be generated through the system).

### 3.5 Examples in practice

This menu of options is valuable in theory, but it's also worth briefly exploring some examples in practice. One of the most notable public positions in this space comes from Google: "We generally seek to share Google research to contribute to growing the wider AI ecosystem. However we do not make it available without first reviewing the potential risks for abuse. Although each review is content specific, key factors that we consider in making this judgment include: risk and scale of benefit vs downside, nature and uniqueness, and mitigation options." [18]

Beyond Google, a number of labs have had to consider these issues. As mentioned earlier, the researchers behind Face2Face and those behind many other synthetic media systems have chosen not to share their code. Some researchers have released code but intentionally made it difficult to use for non-experts.

Different product companies in this space are also exploring mitigations.<sup>6</sup> Synthesia is only working with closely vetted clients. Lyrebird, which enables voice cloning, makes it more difficult to impersonate someone without consent by requiring users to speak particular phrases instead of just training on arbitrary provided data.

## 4 Disagreements around release practices

Different people and groups will have differing views on which kinds of release strategies should be used when. Here we lay out some different dimensions on which people may disagree which affect their views about release strategies for ML research. Our aim is to recognize that genuine divides exist and can lead to polarization of opinion, but that more nuanced discussion can prevent this.

### 4.1 Value trade-offs

Some disagreements stem from fundamental views about the value of openness vs. caution in research.

The ML community has very strong norms around openness: free sharing of data, algorithms, models, and research papers. These strong openness norms appear to be broadly motivated by (1) distributing the benefits of research widely by making it accessible to all of society, and (2) enabling scientific progress by making it easier for researchers to critique and build on one another's work.

Research practices that attempt to limit mal-use by being more cautious about how it is released and distributed necessarily reduce some forms of openness. Some who take openness to be a fundamental value in research may therefore disagree with such practices on principle. However, there are multiple different aspects to openness in research, and, as we've tried to highlight in this paper, multiple different approaches to being cautious about research release. Not all of these will necessarily be in tension with one another, and more exploration of research practices that decrease risk while protecting the most important aspects of openness would be valuable.

### 4.2 Beliefs about risks

Some disagree about the relative size of different risks involved in ML research.

On the one hand, there is the risk that advances in ML might be misused by malicious actors in potentially catastrophic ways, which we've discussed. But restricting the release of ML research also creates its own risks: (1) of increasing power concentration, as a few research groups disproportionately control how ML capabilities evolve, and (2) of creating public confusion or even panic, by creating the impression that advances are more threatening than they are.

Beliefs about the relative size of these risks can lead to two very different perspectives. Those who believe that ML advances will lead to significant harm very soon may want to risk such power concentration in order to safeguard democracy and public trust in the long term. By contrast, for those who think weaponization is less immediately relevant and that we can reassess risks in the future, the costs of restricting research may seem less palatable.

While there is a genuine tension here, it is worth considering approaches that could address both sides of the concern (or at least address one side without exacerbating the other.) For example, some standardization of release practices, potentially managed by external entities, could help mitigate misuse without leading to power concentration.

---

<sup>6</sup>For more on this, see this crowdsourced list of organizations and their actions to mitigate risk: <http://bit.ly/synth-ethics>.



### 4.3 Beliefs about efficacy

Another dimension of disagreement centers not around what the risks are but how effective different practices are likely to be at reducing them.

Given strong incentives or low barriers to develop a technology (or achieve an insight), some suggest it is impossible to prevent either from leading to mal-use the long run, which could mean that restricting the release of research with potential for mal-use is futile. Others suggest that we can significantly impact incentives or barriers, or that slowing down release into the world can still make a significant difference, especially if this gives us time to build defenses against potential mal-use. There is also the perspective that it is easier to build systems to defend against mal-use if more research is public, and the counterview that public information can sometimes help attackers more than defenders ('security through obscurity' may be unnecessary for e.g. keeping data private but is still allegedly crucial for anti-spam defense). As ML researchers continue to experiment with release practices and explore similar challenges in other fields, we may learn about the efficacy of different approaches which can help inform these beliefs.

### 4.4 Beliefs about future needs

Finally, there's a question of whether we might eventually need processes for release of ML research, even if they're not essential now.

For those who believe that we might develop much more advanced ML systems in the relatively near future, and that potential for harm will increase with these advances, then it probably makes sense to start developing careful norms and processes now regardless of current harms. For those who are more skeptical of the possibility of much more advanced capabilities, think that such capabilities are unlikely to be dangerous, and/or that restricting release is unlikely to be effective in the future regardless, developing such processes now looks unnecessary.

\*

Part of the reason for laying out various different options for release of research is to show that this needn't be a polarized debate: it's not a simple choice between 'open' or 'closed' ML research. It's worth considering whether, within our menu of options, there are approaches which can strike a balance between the differing perspectives outlined here.

## 5 Recommendations

We've laid out some considerations, tools, and options for thinking through release of potentially harmful research in a nuanced way. But what must be done now? Here are some brief recommendations:

1. **Increase understanding** of the risk landscape and possible mitigation strategies:
  - *Develop standardized language* for talking about these issues e.g. around hazards, adversaries, mitigations and release options.
  - *Map risks* of different types of ML research in collaboration with subject matter experts, such as e.g. misinformation security researchers for synthetic media. Map out both immediate direct threats and potential longer-term path dependencies, in ways that address researcher concerns around risk hyperbole. Develop practices for safely discussing such risks.<sup>7</sup>
  - *Map mitigation options*, e.g. ways of reducing the harms resulting from mal-use of synthetic media research, and the stages/times at which they are applicable.
2. **Build a community and norms** around competency in understanding the impacts of ML research:
  - *Establish regular workshops* to focus on release challenges.
  - *Spread awareness* of the risks of ML research to both groups who might be affected and who can help mitigate the risks. Proactively seek to include and learn from those who have been impacted.
  - *Encourage impact evaluation* both positive and negative, for research publications, presentations, and proposals (such as that proposed by the ACM FCA [19]).
3. **Fund institutions and systems** to grow and manage research practices in ML, including potentially:

---

<sup>7</sup>This type of investigation might also be referred to as e.g. threat models, risk analysis, or impact analysis, each of which involves a different set of useful lenses.

- *Support expert impact evaluation* of research proposals, so that the burden of this does not fall entirely on individual researchers (who may not have the relevant expertise to assess hazards). This might involve e.g. identifying groups with subject matter expertise who can do evaluations (at the request of researchers), coordinating, and potentially even paying for review.
- *Prototype vetting systems* to help enable shared access to potentially sensitive research (as opposed to the current system where researchers attempt to validate if those requesting their models are malicious actors [20], often via error-prone ad-hoc Googling).
- *Develop release procedures* for research already deemed to raise potential risks (managing all of the above if needed, so that individual researchers can spend more time on actual research while still mitigating risks). Currently, organizations are unilaterally not publicly releasing results, so developing better procedures could actually open up research.

## 6 Conclusion

It is clear that advances in ML have the potential to be misused: the main example we have discussed here is how advances in synthetic media creation may be used to sow disinformation and mistrust (but many others can be, and have been discussed [2]). We must start thinking about how to responsibly safeguard ML research.

Here we focus on the role of release and publication practices in preventing mal-use of ML research. The idea that we might sometimes restrict research release has been met with understandable concern from parts of the ML community for whom openness is an important value. Our aim here has been to decrease polarization in this debate; to emphasize that this is not a simple choice between “open” and “closed” ML research. There are a variety of options for how and when different aspects of research are released, including many drawn from parallels to existing fields, and many possible processes for making these decisions.

There will always be disagreements about the relative risks and benefits of different types of research, the effectiveness of different mitigation strategies, and ultimately how to balance the values of openness vs. caution. We must more deeply explore the risks and options, and develop release strategies and processes that appropriately balance and manage the trade-offs.

Ultimately, we want research to benefit humanity. We see this work as part of a maturing of the ML community, alongside crucial efforts to ensure that ML systems are fair, transparent, and accountable. As ML reshapes our lives, researchers will continue to come to terms with their new powers and impacts on world affairs.

## References

- [1] Jess Whittlestone, Rune Nyrop, Anna Alexandrova, Kanta Dihal, and Stephen Cave. Ethical and Societal Implications of Algorithms, Data, and Artificial Intelligence: a Roadmap for Research. *London: Nuffield Foundation*, 2019.
- [2] Miles Brundage, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, Paul Scharre, Thomas Zeitzoff, Bobby Filar, et al. The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation. *arXiv preprint arXiv:1802.07228*, 2018.
- [3] Ali Breland. The Bizarre and Terrifying Case of the “Deepfake” Video that Helped Bring an African Nation to the Brink. *Mother Jones*, 2019.
- [4] Nic Ker. Is the political aide viral sex video confession real or a Deepfake? *Malay Mail*, 2019.
- [5] Tero Karras, Samuli Laine, and Timo Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019.
- [6] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody Dance Now. *arXiv preprint arXiv:1808.07371*, 2018.
- [7] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2387–2395, 2016.
- [8] RJ Skerry-Ryan, Eric Battenberg, Ying Xiao, Yuxuan Wang, Daisy Stanton, Joel Shor, Ron J Weiss, Rob Clark, and Rif A Saurous. Towards End-to-End Prosody Transfer for Expressive Speech Synthesis with Tacotron. *arXiv preprint arXiv:1803.09047*, 2018.

- [9] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language Models are Unsupervised Multitask Learners. *OpenAI Blog*, 1(8), 2019.
- [10] Drew Harwell. Fake-porn videos are being weaponized to harass and humiliate women: ‘Everybody is a potential target’. *Washington Post*, December 2018.
- [11] BBC News. Fake voices ‘help cyber-crooks steal cash’, July 2019.
- [12] Raphael Satter. Experts: Spy used AI-generated face to connect with targets. *AP News*, June 2019.
- [13] Claire Leibowicz, Steven Adler, and Peter Eckersley. When Is it Appropriate to Publish High-Stakes AI Research? Partnership on AI blog post, April 2019.
- [14] Nick Bostrom. Information Hazards: A Typology of Potential Harms from Knowledge. *Review of Contemporary Philosophy*, 10:44–79, 2011.
- [15] Gregory Goth. Propaganda Bots vs. Transparency. *ACM News*, 2017.
- [16] World Health Organization. *Laboratory Biosafety Manual*. Geneva: World Health Organization, 2004.
- [17] Kyle B Enfield and Jonathon D Truwit. The Purpose, Composition, and Function of an Institutional Review Board: Balancing Priorities. *Respiratory care*, 53(10):1330–1336, 2008.
- [18] Google. Perspectives on Issues in AI Governance, 2019.
- [19] Brent Hecht, L Wilcox, JP Bigham, J Schöning, E Hoque, J Ernst, Y Bisk, L De Russis, L Yarosh, B Anjum, et al. It’s Time to Do Something: Mitigating the Negative Impacts of Computing Through a Change to the Peer Review Process. *ACM Future of Computing Blog*, 2018.
- [20] Rowan Zellers. Why We Released Grover, July 2019.

*We are grateful to the many researchers across academia, industry, and civil society who provided invaluable feedback. This is meant to be just the start of a conversation, and we are likely to update this document as we learn more. If you have thoughts or feedback, or would like to otherwise contribute to this discussion, please reach out to the authors via email or on Twitter.*