

The politics of past and future: synthetic media,  
showing, and telling  
(Published online first in *Philosophical Studies*)

Megan Hyska

Theorists in a wide range of disciplines are now working to come to grips with the promises and disruptive potentials of realistic synthetic media created through generative artificial intelligence. Such media include static images, as well as the dynamic synthetic media we now commonly call “deepfakes”, which are increasingly indistinguishable from genuine photographs and videos, respectively<sup>1</sup>. This paper draws attention to a consequence of such synthetic media: the presence of these media in our information environment contracts our ability to *show* one another things, even as it may increase our resources for *telling* one another things. And it is the project of this paper to draw out the significance of this flux in our communicative powers for the practice of politics. How we adapt to these new technologies implicates longstanding issues of respect, equality, and effability in political communication.

One widely noted use of these synthetic media is, of course, to trick people into thinking that the depicted events really happened when they didn’t. Attempts to use deepfakes in this way are already familiar: in March 2022, a deepfaked video of Ukrainian President Volodymyr Zelenskyy appearing to tell Ukrainian troops to stand down [Simonite, 2022] was circulated online; in May of 2022 a deepfake of Elon Musk promoting a cryptocurrency scam likewise did the rounds [Elon Musk [@elonmusk], 2022]. And there is every reason to imagine that deceptive uses of deepfake technology will continue, perpetrated by everyone from internet scammers to states; the Brookings Institute has recently argued that “democratic governments will almost certainly consider generating and distributing deepfake content” [Byman et al., 2023, 11].

But there is a second use of synthetic media, parasitic on straightforwardly deceptive uses, which is to sow a general skepticism about photography and videography. The philosophical literature on synthetic media has mostly concentrated on this use<sup>2</sup>. As Rini [2020, 7] puts it, “the most important risk is

---

<sup>1</sup>Synthetic audio media is also a topic of significance, but one that this paper leaves aside. For a thorough taxonomy and discussion of both audio and visual synthetic media, see Millière [2022]. While I will often abbreviate the target here to “synthetic media,” as discussed in §2.2, the specific class of interest contains those media that are a) created using deep-learning models and b) difficult or impossible to distinguish from non-synthetic photos and videos.

<sup>2</sup>Exceptions include Rini and Cohen [2022] and Ohman [2020], who note that some of the

not that deepfakes will be believed, but instead that increasingly savvy information consumers will come to reflexively distrust all recordings.” While this articulation suggests the worry that people simply won’t form beliefs on the basis of video, Rini’s full account entails that, even if they did, these beliefs might not constitute knowledge: whereas videos have formerly been sources of perceptual justification for belief, the advent of deepfake technology, Rini says, renders them sources of merely testimonial justification, which is susceptible to defeat by mistrust of the testifier.

Nor is Rini the only one to worry that deepfakes alter the capacity of veridical videography to give rise to knowledge. Drawing on the notion of informativity familiar from Skyrms [2010], Fallis [2021] argues that, in a deepfake-rich environment, videography actually carries less information. And Matthews [2023] argues that, in an environment rich with deepfakes, a veridical video is analogous to a real barn in Goldman’s [1976] false barn country, so beliefs based on it are too lucky to count as knowledge. It seems, indeed, that the many strands of analytic epistemology converge at the conclusion that deepfakes degrade the role that videography will or should play in the formation of belief, generating what Habgood-Coote [2023] has recently termed the “Epistemic Apocalypse narrative” of deepfake technology.

Where the existing literature, then, is deeply focused on how synthetic media will make a difference for would-be knowers—the viewers of photographs and videos—what I want to make visible is the difference this technology will make for would-be *communicators*—those who embed photos and videos in their speech acts. And while some other recent work [e.g. Pierini, 2023, Roberts, 2023] has taken up issues related to synthetic media’s communicative dimensions, it has always ultimately spelled out the disruptive significance of the technology in terms of interruptions to knowledge acquisition. In contrast, the disruptive effect that I point to is characterized in terms of political morality. Ubiquitous synthetic media, I will claim, alter our capacity to use photos and videos to *show* one another things. And I argue that this matters because opting to show rather than tell functions as a way of preserving a kind of relational equality between agents. Alongside their action-additive potentials then, synthetic media also disrupt an era in which the material for the expression of this respect were easy to come by. And this modifies the set of resources available for the harmonization of collectivity with a respect for the individual.

In §1 I review the distinction between showing and telling. In §2 I characterize showing and telling as technological affordances— that is, as things that technologies like videography, photography, and generative AI can make harder or easier. In §3 I go on to say why the rise of technology that makes showing harder matters: showing is a way that we preserve relational equality through superficial communicative asymmetries and express respect for our audience. This is particularly significant, I argue, when we are establishing a shared understanding of the past that collective political action is to be predicated upon.

---

earliest, and still most prevalent, uses of deepfake technology are pornographic. In these cases, the primary goals may be to titillate or to humiliate, but neither of these necessarily require that anyone be genuinely deceived about whether what they’re watching is real.

And in §4 I close by considering the way that synthetic media created through generative AI gives us new ways of *telling* one another things that are important for political collectivity: our visions, both hopeful and horrified, of possible futures.

## 1 On showing and telling

In a well-known series of papers, H.P. Grice [1957, 1989] draws a distinction between “‘deliberately and openly letting someone know’ and ‘telling’” [1989, 218]. This distinction has subsequently often been rendered as that between *showing* and telling. Cases of telling, for Grice, involve three components<sup>3</sup>:

*Informative intention:* the signaler intends to bring about a particular cognitive response in her audience.

*Communicative intention:* the signaler intends that her audience see that she has the informative intention.

*Causal Intention:* the signaler intends that her audience satisfy her informative intention at least in part because they satisfy her communicative intention.

These three conditions are satisfied in many paradigmatic cases of linguistic communication; if I say “aloe is good for sunburns,” I do so because I want to change your credence in this proposition, or perhaps just make it more salient to you—informative intention, check. What is my plan such that I expect making certain noises in your vicinity to bring about that alteration in your cognitive environment? I expect you will take my noises as an indication that I want you to have this particular cognitive response (communicative intention, check) and that your recognition that I want you to have this response will get you to in fact have it (causal intention, check). An interesting thing about telling then is that it functions on the premise that the audience will be moved to by a recognition of the signaler’s communicative desires to in fact alter their minds. It presumes a trust in, or dependence upon, the speaker.

Contrast telling now with showing. In showing cases, the signaler still has informative and communicative intentions— they still want to touch the audience’s mind somehow and for the audience to see that they want this— but they don’t plan to bring about the desired cognitive response *via* the audience’s recognition of their intention to do so. Rather, in an act of showing, one intends that the desired cognitive response be caused by the audience’s recognition of some independent evidence. So showing, as we will characterize it, involves something like the negation of what we above called the causal intention: in cases of showing, the signaler intends that the recognition of their informative intention be preempted, as a cause of the desired cognitive change, by the other

---

<sup>3</sup>This presentation, while still basically faithful to Grice’s vision, is not quite Grice’s own; it borrows from post-Gricean commenters including Strawson [1964], Sperber and Wilson [1986], Neale [1992].

evidence they make available<sup>4</sup>. In one of Grice’s examples [1957, 382], King Herod presents Salome with the head of Saint John the Baptist— he intends her to come to believe that St John the Baptist is dead (informative intention, check), and he intends her to see that he intends this (communicative intention, check), but he relies on the independent evidence, the fact that St John the Baptist’s head is detached from his body, to convince her that he is dead.

Note that it follows from the continued necessity of the informative and communicative intentions that an act doesn’t count as one of showing, in our sense, just because it *happens* to bring about cognitive alteration via some environmental evidence. If I just so happen to move out of your line of sight en route to go get a glass of water and as a result you see, and come to believe, that there is a tree out the window which my body was previously obscuring, I haven’t *shown* you this. What is more, an act doesn’t count as showing if the signaler intends to change the audience’s cognitive environment but doesn’t want the audience to know this— so even supposing that, when I go to get my water I also *want* you to notice the tree, this still doesn’t count, unless I intend that you infer this intention. In other words, showing, like telling, still has to be *ostensive*, i.e. it has to “[make] manifest an intention to make something manifest” [Sperber and Wilson, 2015, 49].

In cases of successful showing, then, the audience does still recognize the signaler’s informative intention, but the causal path that terminates in the cognitive response the signaler hoped for doesn’t route through this recognition. One interesting puzzle is why, in acts of showing, we should care whether our audiences recognize our informative intention if their doing so isn’t instrumental in these intentions’ satisfaction<sup>5</sup>. That we *do* care reveals a concern with more than the impersonal orchestration of the other’s attitudes. We are concerned with an explicitly *inter*-personal state of affairs, not one confined to the interior of the other’s mind. We will be in a position to return to this puzzle with some novel solutions in §3, when we consider the ethical significance of showing in political contexts.

First, however, we investigate how both showing and telling can be thought of as technological affordances; that is, as acts made easier by some technologies, and harder by others.

## 1.1 Multimodal showing and telling

Acts of communication are often multimodal. If, in the act of describing your brother’s red hair you decide to draw me a picture of him, then the primary medium of your speech act may be spoken language, but also involved now is

<sup>4</sup>This is the one part of my presentation of showing which I take to be a bit non-standard; often, showing is characterized as involving the mere *absence* of the causal intention, rather than by the affirmative presence of an intention to the contrary. I think however that my addition here still captures most if not all of the cases of showing that theorists have been interested in discussing, and that it sharpens the showing/ telling distinction in a way that is helpful for our subsequent discussion.

<sup>5</sup>Hyska [2023b] addresses a related puzzle about the value of intention recognition to communication generally.

ink on paper. And if, in the act of saying how beautiful the sunset is, I find that words don't suffice and I lead you to the window to see it yourself, then the world outside the window is now recruited as one medium by which I communicate to you.

Multimodality appears in both showing and telling: in the above example, the drawing functions as a form of evidence that your brother's hair is red only insofar as it is evidence that you intend me to believe that his hair is red. Because it is on the basis of my recognizing that informative intention that you plan for me to come to satisfy it, this multimodal speech act is an example of telling. But in the example where I lead you to the window to see the sunset, I have a plan for getting you to believe that, say, the sky is the wildest gradient of red to pink to orange, which doesn't require you to believe it on the basis of inferring that I intend you to do so, but on the basis of some independent environmental evidence. So this is an instance of showing.

While parts of the natural world (e.g. trees, severed heads, sunsets) can be embedded in speech acts, the above example involving drawing makes clear that human artefacts can also be so embedded. It might be tempting to suppose that speech acts embedding artefacts will always be instances of telling, but in fact it isn't so. As defined, showing is about the ability to provide justification that is independent of the inferred informative intention, where a source of justification, A, is independent of another source, B, just in case the defeat of B would leave the justificatory force of A untouched. And some artefact-types are such that their tokens have epistemic justificatory force that is independent of their wielders' inferred intentions in this way. This is because, while an artefact as such is created by human intervention, some artefact-types are such that even a skilled human's desire is very unlikely to suffice to bring about a particular desired token.

Consider something like a block in a block-chain, which contains transaction data that is linked, by a cryptographic process, to the data in the previous block in the chain. While obviously blocks are artefacts, the promise of this technology is precisely that, because the ledger is widely distributed, no one person can falsify transaction data by fabricating a block that isn't actually in the chain. When I present you with some part of a blockchain transaction ledger then, this is quite different from providing you with a hand-drawn picture. The drawing was entirely under my power to create—its creation was “directly mediated by the producer's desires and beliefs” [Millière, 2022, 3]—and if the use that I am making of my ostensive powers comes into question, then the drawing, being either definitely or at least plausibly an expression of these powers, likewise falls under suspicion. The block chain however, is not in my power to fabricate; so even if you don't trust me to use my ostensive powers honestly, the blockchain, falling outside these powers, provides justification for a certain transaction record. The justificatory independence of the blockchain means that speech acts embedding it can show, not merely tell, that for instance a currency transfer went through.

An important question is whether the sort of independence that makes a multimodal speech act one of showing has to do with what could *in fact* have

been fabricated by the signaler, or what the audience *takes it* could have been fabricated. There is a certain ideal epistemological project that might lean in the former direction; however, that is a different project from the one that Grice, and we, have in mind. Telling, for Grice, is characterized by the fact that “the recognition [of the informative intention] is intended by [the speaker] to play its part in inducing the belief, and if it does not do so something will have gone wrong” [Grice, 1957, 383]. In other words, telling is characterized by the mechanism by which the speaker intends to actually bring the response about in the audience. And for Grice, as for many action theorists, this intention is itself constrained by what the speaker believes to be possible—i.e. by what they think the audience’s mind is actually like<sup>6</sup>. So, in asking whether a medium enables showing or else mere telling, the question is not whether the medium really *is* beyond the powers of a skilled individual to fabricate, but whether people take it to be so.

## 2 Showing with photos, videos, and synthetic media

This brings us to the question of whether photography and videography have been regarded as speaker-independent in a way that makes them fit for embedding in acts of showing. It’s worth noting up front that whether the embedding of photography and videography in a speech act can make for an instance of showing is distinct from the question, much discussed elsewhere [see e.g. Walton, 1984, Cavendon-Taylor, 2013, Rini, 2020], of whether these media provide perceptual, as opposed to testimonial or inferential, justification for beliefs. The showing/ telling distinction turns on the independence of the proffered evidence from the speaker’s intentions, not on the flavor of justification that this evidence offers. Indeed, the two things come apart: acts of showing needn’t embed modalities that provide perceptual justification at all. They can also embed evidence that provides mere inferential justification (think of showing someone your watch), and probably also testimonial evidence (consider gesturing at another speaker who is making a speech act).

We will first consider the character of photography and videography in an era prior to generative AI. After this, we will be in a position to assess how the availability of synthetic media modify the affordances of these earlier technologies.

---

<sup>6</sup>The choice to construe the telling/ showing distinction in this way for the purposes of this paper is rationalized not just by what Grice intended, but also by the fact that, in §3, I’ll ultimately be concerned with the ethical significance of signalers being able, or not, to rationally intend to communicate in certain ways. So it makes sense to draw a distinction that tracks what speakers rationally intend rather than what they *would* be able to rationally intend if they took their audiences to have all the technological facts.

## 2.1 Photography and Videography

§1 put us in a position to see that whether the embedding of photos and videos can make for acts of showing turns on whether photographs and videos are artefacts that people commonly believe others to be able and willing to fabricate at will. Now, the technical capacity for the modification of photographs—through retouching of both negatives and prints, multiple exposures of a single negative, or creating a single print with elements of multiple negatives—has been available since the 19th century [Fineman, 2012]. And the availability of digital photography and associated digital editing software to consumers starting in the 1990s of course made this manipulation all the easier.

However, it is important to bear in mind that our willingness to believe that a medium is the product of fabrication rather than of an appropriately speaker-independent documentary process is a function not just of technical limitations, but also of economic and social ones. Professional documentary photography, at least, is enshrined in norms that reward technical competence and honesty, and punish all but a small set of photographic manipulations (e.g. cropping, minor color re-balancing) [Cohen and Meskin, 2004, Walden, 2008, Abell, 2010, Hopkins, 2012, Habgood-Coote, 2023]. And our awareness of these norms seems to have gone some way toward assuring us that many of the photographs we see aren't fakes, even though technical limitations alone haven't prevented this.

An extreme social constructionist version of this position would say there is in fact *nothing* about photography that gives it a better claim to independence than any other medium *except* for the social norms that have, as a contingent matter, come to govern its use. Lopes [2016, 112] seems to endorse something like this position when he argues that “Epistemically virtuous images come not from photography but from belief-independent feature-tracking. Standard photography secures its virtue by regimenting the design and use of technology to ensure belief-independent feature-tracking.” On such a view, hand-drawn stick figures could come to be treated as speaker-independent in just the way photographs have often been, if only we had a different set of professional norms around them.

Rini [2020], while still regarding norms as an important determinant in the epistemic role that a technology is able to play, seems to reject this degree of social construction. Per Rini's account, a key reason that photography has been regarded as epistemically independent of the communicator, is because of the technical properties of another technology: videography. While photography has long been technically fakeable in all sorts of ways, serious, convincing videographic modification has up until recently required significant time, money, and expertise. A 2005 *Popular Science* article fretting about the duplicity enabled by photoshop noted that, “One bright spot is that for now, at least, we only have to worry about still images,” because, “there 's no Photoshop for movies” [Casimiro, 2005]. Rini [2020, 13] has argued that the relative non-manipulability of video has enabled video to “acutely correct and passively regulate photographic evidence”; in other words, a fabricated photo can be outed by a video of the same scene, and expectation of this possibility functions as a deterrent

to would-be photo fakers. This account clearly still makes crucial use of norms, since the existence of a technology that could uncover photographic fakes would be irrelevant unless there were norms to make the prospect of such unmasking practically important. But, unlike on a pure social constructionist account, norms do this work alongside the intrinsic properties of available technologies.

We see then that there are a variety of accounts of how photography has, notwithstanding its technical susceptibility to manipulation, continued to be treated as beyond the powers of the individual to fabricate. This bears out the claim that speech acts embedding photography can count as acts of showing in a way that hand-drawn images cannot. And videography, the faking of which has been mainly ruled out merely by technical considerations, has had all the more claim to being a method of showing.

What we can now appreciate is how the invention of these technologies made a difference to our communicative options. While showing had always of course been possible, photography, and later videography, enabled acts of showing that hadn't been available before.

Consider for instance the first photographs of objects in motion, achieved by Eadward Muybridge in the 1870s. Before this time, long exposure times had made it impossible to photographically capture a moving object without blurring. But in 1877, Muybridge's innovations with photographic emulsions, shutter speeds, and trip wires allowed him to capture a series of pictures of a horse running at a race track in Sacramento.

And what the camera captured was, to some, unexpected. Muybridge biographer Rebecca Solnit recounts what happened when Muybridge's patron, Leland Stanford, presented the photographs to the realist painter, Jean-Louis-Ernest Meissonier:

Meissonier at first thought that the pictures had been edited, eliminating one position of the gallop, the position he had often painted. When he realized the sequence was complete, he declared, "All these years my eyes have deceived me." Stanford reportedly replied, "The machine cannot lie," and gave him a tutorial on horses' gaits, which ended with the painter exclaiming in despair "After thirty years of absorbing and concentrated study, I find I have been wrong. Never again shall I touch a brush!" [Solnit, 2003, 197]

Muybridge's photographs showed what had previously been unshowable due to limits on the human ability to discriminate all the stages of a rapid movement. But this was far from the only way in which photography and videography augmented our capacities for showing. Indeed, the most obvious way in which they did so was by allowing events that were directly observable only at some point in the past to be shown in the present; or events that were only directly observable in one geographical location to be shown in another. With photography, "the annihilation of time and space" upon which 19th century industry was everywhere intent [Solnit, 2003, 11], extended advantages to the activity of showing.



## 2.2 Synthetic media and the revision of old affordances

As discussed in the last section, the ability to manipulate photographs, at least, is not particularly new. What is worth noting though is that the familiar varieties of photographic manipulation tend to involve the alteration of a pre-existing photograph. Millière [2022] offers a taxonomy of audio-visual media which distinguishes between *totally* synthetic and merely *partially* synthetic cases. Partially synthetic media include both things like a digital photo after the application of an Instagram filter (what Millière calls *partial global* alteration) or the editing out of a figure in the background (*partial local* alteration).

What is novel about the next generation of deep learning (DL)-based synthetic media is that it can be both *totally* synthetic, and generated by users with very little expertise. Now-familiar text to image models (e.g. those behind DALL-E, Midjourney, or Stable Diffusion) are trained on huge corpora of real photographs (as well as non-photographic images), but allow users to generate wholly new quasi-photographic media with none of the skill or time-investment historically required by even very good digital editing software.

At the time of this writing, the most widely accessible use of DL models to generate dynamic synthetic media (i.e. media that look like videos) involves the creation of merely partial local synthetic media. These outputs are based on existing video and still image sources, but modified through face swapping, head puppetry, or lip syncing techniques [Tolosana et al., 2020, Zakharov et al., 2019, Prajwal et al., 2020]. Some rudimentary tools for the creation of totally synthetic video based on a text prompt are now available to the public, but with outputs that are more funny than convincingly videographic. But it is a widespread expectation that tools for the creation of convincing, totally synthetic audio-visual samples are on their way.

The DL-based synthetic media we see today are, and will only become more, convincing. For a while now, the best DL-based image generators have been able to produce images of human faces that real humans can detect at rates no better than mere chance [Hulzebosch et al., 2020, Lago et al., 2022, Nightingale and Farid, 2022]. And while certain artifacts and distortions have worked as tells historically (e.g. the misrendering of human hands), there is evidence that as the models get better, these tells are disappearing [Edwards, 2023]. Finally, DL-generated media are often uniquely equipped to evade technological detection. Many of these models are trained via generative adversarial networks (GANs) in which a generator node iteratively works to produce an image that won't be caught by a discriminator node, and generates better and better results over time as a result of the discriminator's feedback. This adversarial architecture positions these networks to co-opt any detection software for the improvement of their own discriminator nodes, which will then train the next generation of models to be undetectable [Farid, 2022].

Insofar as these synthetic media become increasingly accessible, their production is similar in kind to the creation of a simple pen-to-paper drawing: nearly anyone can do it, to aid in their communicative purposes whenever they please. As a result, when an audience *knows* that they are being presented with

DL-based synthetic media, the act of presenting it to them can at most be an act of telling. But because these synthetic media are so realistic, even when people are presented with real photos and videos, they won't know that this is so. As Pierini [2023] puts it, this places people in a position where they don't know whether the media on offer features an "intentional" or "non-intentional" standard of correctness—as we might put it, the audience doesn't know whether the evidence provided is properly treated as dependent or independent of the speaker.

The capacity to perform an act of showing, as we have characterized it, turns on whether the signaler believes that the modality she embeds in her speech act is, independent of the audience's recognition of her intentions, sufficient to bring about the cognitive response she desire to see in her audience. Can would-be showers believe this, of photographs and videos, in an environment rich with DL-based synthetic media? Consider first a case where the audience does not trust the signaler and the signaler knows it; the signaler then cannot rationally believe that the audience will take the proffered media as a photo at all, and so cannot believe that it suffices to bring about the desired cognitive reaction in them. The signaler therefore cannot engage in an act of showing through the use of photos. Now consider a case where the signaler knows that the audience does trust her. This signaler must anticipate that, upon presenting the audience with a photograph, the audience will reason that it is a photograph because they trust her. While the medium of photography as such might still be said to offer speaker-independent justification, the justificatory force of any token photograph is now dependent on the signaler's inferred intention to be honest. The same goes, *mutatis mutandis*, for videos. So the embedding of these media in a communicative act now counts merely as a means of telling, not of showing. The charge then is that our new capacities for synthetic media effectively modify the communicative affordances associated with the older technology of photography and videography.

An immediate objection to this claim follows from what we have said above about how a medium can continue to enable showing even when technical considerations bring it under the powers of a signaler to fabricate. If photography has, through the 19th, 20th and early 21st centuries, functioned to enable showing because of an appropriate set of surrounding norms, why not suppose that the same will be possible in an age of DL-based synthetic media? This transposes the worry that Habgood-Coote [2023] has raised for the existing "Epistemic Apocalypse" literature over to the case I have made here. Rini's [2020] account anticipates this objection, and responds that norms that punish photographic tampering have, for a long time, worked alongside the technological "backstop" of videography. After all, a norm that punishes a certain action retains its acute and passive regulatory functions in proportion with how possible it is to catch those who've committed that action. On Rini's account then, dynamic DL-based synthetic media (i.e. deepfakes) change the game because they get rid of this backstop, and no set of norms that penalize faking can function well without one.

Habgood-Coote [2023, 4] charges that Rini's position relies, fallaciously, on

the idea that videography has been specially equipped to be a backstop because it hasn't itself relied on norms governing its practice; applying ideas from Goldberg [2020], he argues that in fact *all* technology relies diffusely on a community of practitioners who enforce its associated norms. In this case, we've been functioning without a backstop all along, and if no crisis has befallen us so far, there's no basis to forecast one now. But this does not seem to me to touch the essence of Rini's suggestion. Were all social norms suspended in 2005, videography would still have been much harder to fake than photography—observing this much does not require denying that social norms were in fact in place as an extra guarantor of videography's veridicality. We can acknowledge that the intrinsic features of a technology play a role in enabling and constraining the social dynamics that play out around it without swinging all the way toward a technodeterminism that refuses to acknowledge any role of the social in informing the uses of the technology. Indeed, the notion of a technological “affordance”—which Habgood-Coote avails himself of, just as do I—has often been used in technology studies specifically to express this middle position between pure social constructivism and hard technodeterminism [Hutchby, 2001]. Since the coining of the notion in Gibson's [1979] work on perception and its uptake by the literature on artifacts in Norman [1988], to speak of an object as having affordances has been to grant that it has intrinsic features which “frame, while not determining, the possibilities for agentic action” [Hutchby, 2001, 444]<sup>7</sup>.

Where this leaves us, I think, is with a recognition that the intrinsic features of video historically worked alongside a system of norms to allow it to play a special role in the technosocial ecosystem. It also suggests that, holding fixed current social norms, the intrinsic features of new synthetic media alter videography's ability to play this role. As a result, the communicative possibilities open to us have changed. This is all entirely compatible with the proposal that, instead of forecasting certain doom, those of us writing about emergent synthetic media technologies should be “thinking about how to design norms for techno-social practices” Habgood-Coote [2023, 20]. Far from denying the possibility of adaptation to this new technology, I take the project here to be developing a problematic that helps us envision and assess the significance of possible adaptations. For instance, there are a number of ideas circulating for how blockchain technology could effectively come to function as a new backstop for both photography and videography. On some suggestions, blockchain is used to tether images to their metadata so that it's easy to track their provenance across platforms [Koren, 2020]<sup>8</sup>. And Chesney and Citron [2018, 1814] suggest that the ubiquity of deepfakes could incentivize blockchain-backed location logs for anyone worried about being able to conclusively refute a deepfake that represented them as doing something embarrassing or criminal in another loca-

<sup>7</sup>For a more extensive history of the notion of an affordance, see Davis [2020, 25–43].

<sup>8</sup>This proposal has originally been associated with solving the problem of lost context or else of photoshop-era manipulation, rather than DL-based synthetic media. However, it has clear applications here as well, at least in determining whether the image traces to a source you trust not to have fabricated it.

tion. These adaptations hold *prima facie* promise for addressing some (though I think not all) of the disruptions that will foreseeably come with DL-based synthetic media, though with significant tradeoffs (e.g. in the location log proposal, constant self-surveillance). But without articulating precisely what these disruptions are, we would not be able to assess any proposed adaptation’s merits.

So, I don’t take it that the purpose of this paper is to describe a terminal predicament which no technosocial adaptation could possibly navigate, but to articulate one disruptive potential that new social and technological arrangements will arise to mediate. What I go on to do now is illustrate why the shift that DL-based synthetic media will bring about in our communicative powers *matters*. Rather than trying to exhaustively document the ways in which this shift might be significant in all domains of life, I look closely at one particular domain: politics.

### 3 Showing and the Ethics of Communication

As we saw earlier with the example of Muybridge’s horse, the enlarged powers of showing that we acquired through the technologies of photography and later videography were hugely consequential for many fields, including art and science. Documentary photography and videography<sup>9</sup> also became a widespread tool for politics.

Some of the most vivid examples of political showing through photo- and video-graphy involve the documentation of scenes or events which communicators present as either very good or very bad. In the 1920s and 30s, countries around the world saw an organized movement of working class people bringing cameras into their workplaces and neighborhoods to depict both “the beauty of labor itself,” the “horrors of social misery” (12), and the victories of the labor movement, in order to challenge the narrative of bourgeois photojournalism [Ribalta, 2011, 12]. During the Great Depression, the US Farm Security Administration produced hundreds of thousands of photographs of “agrarian scenes of human distress” [Corbus Bezner, 1999, 6] with the goal of generating the political will to enact change. In the 1940s, Gordon Parks photographed the realities of life in Harlem to challenge the conditions in which, as his collaborator Ralph Ellison put it, Black Americans were “the displaced persons of American democracy” [Parks, 2016]. And in the contemporary United States [Richardson, 2020], Egypt [Tufekci, 2017, 22-27], and elsewhere, photos and videos capturing police brutality and its aftermath are responsible for triggering the most massive popular uprisings of recent memory.

---

<sup>9</sup>I have in mind a rough distinction between those photos and videos that were presented as first-hand evidence of actual events, and those that weren’t. The limits of the “documentary” are of course contested [see e.g. Corbus Bezner, 1999, 1-15]. I am conceiving of the category here as tolerating a certain amount of staging or posing. In any case, our emphasis on documentary media excludes many famous political propaganda films. These are of course fascinating in their own right, but their relationship to the act of showing is more complicated than documentary media’s.

Of course, photography and videography have been used by actors of every political stripe. In Fascist Italy, Mussolini lauded the ability of the photography to “show the world what the Blackshirts have achieved in all areas of activity” [Mussolini, 1932]. And in the United States, the mutilated bodies of Black lynching victims were, through the 19th century and well into the 20th, routinely photographed surrounded by the grinning mobs that had murdered them—indeed, it was at the behest of the mob that the photos were taken. These photos were then passed around as postcards, and even publicly exhibited in moving picture theaters, for decades after the fact. While the NAACP would later use these same photos to show the moral depravity of lynching, they were originally purported to show the moral uprightness of the mob, and deployed on behalf of the political projects of white supremacy [Wood, 2009, Medina, 2018].

As different as these cases are, what they have in common is a desire to establish a consensus about some feature of the past, whether recent or distant: that it has been bad, and the future must, at all costs, not resemble it; that it was good and must be repeated; that it holds evidence of our vulnerability; that it holds evidence of our strength. These claims about the past are then marshalled as the premises for future collective political action. In all these cases, we should ask: why did political communicators aim to show, rather than tell, their audiences about the pasts on which they predicated their politics?

I do not assume that this question has a single answer; that the media we use to show draw more attention and are more emotionally affecting is no doubt one consideration. And in light of the points made by the existing epistemological literature on photography, it is evident that the persuasive efficacy of showing over telling is a further one. Perhaps in some cases, these are the full story. But I want to articulate another reason why a political actor might be interested in showing rather than telling, and so another reason why they might historically have reached for photography or videography as technologies that enable this. In so doing, I hope I can make visible an unexpected thing that we stand to lose as a result of synthetic media.

### 3.1 Showing and Relational Equality

To get at this additional rationale for showing in politics, we begin by considering a case that does not in fact involve photos or videos. In the 19th century, European anarcho-socialists became preoccupied with the tactic they called “propaganda by the deed”. As opposed to propaganda that functioned through pamphlets, posters and speeches, propaganda by the deed consisted typically in direct action—industrial sabotage, destruction of tax records, assassinations of political or industrial leaders. While these actions of course had immediate material consequences, they were conceived of as having also a *communicative* dimension [Cahm, 1989]. And whereas speeches, posters, and pamphlets merely *told* the European working and peasant classes that the ruling classes were wicked and that rebellion was possible, propaganda by the deed aimed to communicate by *showing* them. Why aim to show rather than tell? As the Italian Federation of anarcho-socialists put it in an 1876 letter to their

comrades elsewhere on the continent, “the insurrectional act which is intended to affirm socialist principles by deeds, is the most effective means of propaganda and the only one which, *without deceiving and corrupting the masses*, can penetrate down to the deepest levels of society” [Guillaume, 1910, 116, italics mine]<sup>10</sup>. In other words, propaganda by the deed was motivated not just by a concern with effective persuasion, but also with ethics: with not abusing their audience in the very act of persuading them.

The writers of this letter were not concerned merely with getting others to believe them, but in doing so ethically. Like them, in attending to the importance of showing in politics, I am not concerned with whether a communicative act can bring about knowledge. I contend that even in cases where an audience could be persuaded by either telling *or* showing, and where either would result in knowledge, there is still something lost when showing ceases to be an option. Specifically, I claim that that acts of showing are a way of preserving relational equality through asymmetries in communication, and that as a result, they express a distinctive sort of respect for the audience. The loss of the ability to show then means the loss of the ability to express this respect. And this has moral significance.

To be clear at the outset, I am not claiming that showing is always, or indeed ever, obligatory. I do however think that a communicative act’s being one of showing can add something of moral worth. But while I will argue that showing is a pro tanto morally positive feature of political communication, it does not follow that there is anything wrong with telling. There is no tension between the view I argue for here and the well-developed idea in social epistemology that it can be unjust not to believe someone when they tell you something, for instance [see e.g. Hill Collins, 1990, Fricker, 2007, Medina, 2013]. Finally, I am evidently focusing on the stakes of showing specifically in one domain, collective political action, where my instinct is that relational equality has a distinctive importance. The full scope of the domain in which showing might have this importance is something I leave unspecified here<sup>11</sup>.

The first part of my position to argue for is that showing does something to preserve relational equality. Consider a case: I hope that you will join me in backing a challenger to a corrupt local politician. I come to you and tell you should join the challenger’s campaign because the incumbent is guilty of some serious criminal behavior. I know this because I’ve seen key emails between members of his campaign, and a surreptitiously taken video of his misdeeds. Suppose, moreover, that you’re willing to campaign against the incumbent just based on my telling you that he’s corrupt—you trust me. But suppose that this course of action is not without risk for you; you own a small business and know that the incumbent won’t grant you a liquor license or necessary construction

---

<sup>10</sup>For further discussion of this case and other applications of the showing/ telling distinction to political propaganda, see Hyska [2023a].

<sup>11</sup>Here one is reminded of the analogous debate of the scope of application of the norms of public reason, with some [Rawls, 2001] claiming that it applies narrowly to the hammering out of constitutional frameworks and matters of basic justice, but others [e.g. Larmore, 1996, Gaus, 2011] giving it larger scope.

permits if you side against him. If you and I proceed in our collective action against the incumbent on these grounds, there is an important kind of asymmetry between us: your rationale for action relies on my testimony whereas my rationale for action does not rely on yours. And imagine that in the course of the campaign, more and more relevant information becomes available to me via independent evidence, and available to you only through my testimony. I have a role in shaping your exercise of political agency that you do not have in shaping mine.

Is this asymmetry troubling? Perhaps not, so long as I don't abuse it—so long, that is, as I tell you the truth, don't withhold important information from you, and don't use my power to coerce you into doing anything you wouldn't were the asymmetry corrected. But there is a long tradition of thinking that some power asymmetries are politically problematic even when they are never deliberately leveraged. This tradition finds expression, for instance, in relational (as opposed to distributive) approaches to political equality [e.g. Young, 1990, Wolff, 1998, Anderson, 1999, Scheffler, 2003, Schemmel, 2012]. As Pettit [2001, 137] has put it, if others have an ability to arbitrarily interfere in your life and you don't have a reciprocal ability to arbitrarily interfere in theirs, then even if they don't use this ability, you live “at the mercy of others” , escaping poor treatment only by “the grace or favour of the powerful”. Under these conditions, even if you are free from actual malicious or arbitrary interference, you still stand in an unequal relation to others, and are still, as Pettit would put it, subject to domination.

It is clear that a person might perpetuate relational inequality through the *content* of their communicative act. The slightly different point being made so far though is that the mere fact of communicating, regardless of the communicated content, challenges relational equality when it brings about an asymmetry in epistemic dependence between parties engaged in a collective political undertaking. Keith Raymond Harris [2023, 3] makes this point, noting that a distinctive kind of domination comes about when there is a “significant imbalance in the control each party has over the evidence available to the other.” And Harris thinks that control over what evidence is available to others is often exercised through communication. However, he doesn't appear to think that there are any important qualitative distinctions among types of communication that might modulate these inequalities.

But I think there are. I think showing gives us a way to manage relational equality when communicating that telling doesn't. While acts of showing and telling both represent one person's ability to affect the practical deliberations of another, and while both may involve one party's dependence on the other for some information, in acts of telling that relation of dependence is preserved, whereas in acts of showing it is in an important sense discharged. If I show you something, although I am causally implicated in bringing about your belief, your trust in me is irrelevant to whether you believe it. My communicative act is the scaffold that you climbed up to get there, but your dependence on me terminates once the communication is done. In contrast, in acts of telling, you depend on me not only causally, to bring about the belief in the first place, but

you *go on* depending on me for the justification of your belief<sup>12</sup>.

Sometimes in collective projects, one person will have more project-relevant information than another. In these cases, it is unavoidable that one person will bring about relevant knowledge in the other at a level that is significantly greater than what the other can bring about in them; there will be an asymmetry in knowledge sourcing. The solution to this asymmetry, certainly, is not for the more knowledgeable person to simply shut up and not share what they know; collective political action without shared knowledge tends to be possible only with the help of even more obvious forms of relational inequality. And indeed, many have argued that under certain circumstances we have a moral obligation to share what we know [e.g. Lackey, 2020, Watson, 2022]. Instead, what I have argued is that showing makes asymmetries in knowledge sourcing tolerable because it detaches them from the meaningful relational inequality of asymmetric justificatory dependence.

What I want to add now is the point that, because of this equalizing capacity of showing, acts of showing function also as expressions of respect. Recall the puzzle about showing that I gestured at in §1; why would anyone care about having their informative intentions recognized if, as is definitionally the case in acts of showing, this recognition wasn't instrumental in these intentions' satisfaction? Call this the ostension puzzle. One natural answer to this puzzle is that, in addition to caring about getting some information across to our audiences, we also care about that information becoming something like *common* knowledge or belief<sup>13</sup>. This requires that our audiences see that we already believe just what they have come to believe, which is something they can infer from the recognition of our intentions. This is a perfectly adequate response to the ostension puzzle. However, it has nothing to say about why, if common belief is the goal, the signaler wouldn't just advert to telling, which can achieve this just as well. Of course, in some cases, mere telling will not suffice to bring about even first-order belief in the audience, because they don't trust the signaler—and this explains why one would choose showing as the means to reach common belief. I don't doubt that something like this is right in many cases. But the account of showing's ethical significance developed above suggests a response to the ostension puzzle which simultaneously explains why one would advert to showing rather than telling even in cases where showing and telling are on a par with respect to their persuasive effects. When I show you something, I am

---

<sup>12</sup>Precisely how to spell out this dependence will depend on your position within the epistemology of testimony: for some theorists, the justification we have for testimonial belief always includes the trustworthiness of the speaker; for others, testimony has a default justification and the (un-)trustworthiness of the speaker becomes relevant only as a possible defeater; for yet others, testimony doesn't rely for its warrant on evidence that the speaker is trustworthy, but instead functions as an invitation to treat the speaker as trustworthy, which confers a kind of non-evidential epistemic warrant. I think that each of these positions allows for an acceptable way of spelling out what I have here called dependence.

<sup>13</sup>I am spelling this possible solution out concerning belief for simplicity's sake. As mentioned before, the Gricean tradition has, since Grice [1957], come to say that belief is only one among many kinds of cognitive reactions that an act of ostensive communication might bring about. This solution can be spelled out in terms of a commonly salient set of propositions too, with a few tweaks.



working to preserve relational equality between us; this much might be achieved without you recognizing that I had any informative intention at all. But perhaps I want you to *see* that I am working to preserve relational equality between us. This cannot be achieved unless you recognize my informative intention, but it also cannot be achieved via telling. Only showing will do. Showing is a unique mechanism for the development of solidarity because even as it provides evidence for some particular proposition relevant to the matter at hand, it displays evidence of a disposition to preserve a certain kind of relationship<sup>14</sup> A practical political significance of showing is then that it is an overture to the sort of relationships that make collective action sustainable.

In summary then, showing is a way of preserving relational equality even when one person is, asymmetrically, the source of information, because it avoids an asymmetry of justificatory dependence. And the ostensive character of showing means that it also expresses to the audience one's commitment to this relational equality. It is then an ethically expressive act. If, as I argued in §2, the advent of synthetic media rolls back our capacities for showing, this then constrains the possibility of the ethically expressive act I've just described. Just as we had many ways of showing one another things prior to photo- and videography, we will retain many through the development of increasingly sophisticated DL-based synthetic media. But in those situations where photos and videos have been our only means of reaching across time and space to show one another distant states of affairs, synthetic media threaten to leave us with no way of managing equality through asymmetry

## 4 Coda: on telling and the future

If the affordances of synthetic media are such as to remove some of our communicative options, they are also such as to add others. In particular, new synthetic media represent a new way to *tell* one another things.

An interesting use of deep learning tools like Dall-E, Midjourney and Stable Diffusion is to depict dreams and products of imagination; private, powerful sensory experiences that resist capture in language. In February 2023, students at an MIT hackathon used AI tools to create short videos responding to the prompt “Tell me your dream” [Zhang, 2023]. On subreddits dedicated to witchcraft, users describe these generators as tools of “technomancy,” usable for the purposes of “manifestation”—a mystical act of telling the universe what you want your future to hold.

The use of AI-generated images and videos to tell others about your vision

---

<sup>14</sup>I will note in passing that, while I think the connection of relational equality with solidarity is intuitive and uncontroversial enough, some accounts of solidarity are not obviously constructed with ideals of relational equality in mind. Kolers [2016] for instance conceives of solidarity as necessarily asymmetrical, involving one party deferring to another. It might be argued that, in the cases where Kolers thinks such deference is warranted (typically, when engaging in political collectivity with people who have historically had less power than you), deference actually restores something like relational equality. The proper political role of deference is, however, a matter of substantive debate [see e.g. Táíwò, 2022].

of the future appears also in politics. In March 2023, conservative pundit Jack Posobiec tweeted out an excerpt from his show which included an AI deepfake of US President Joe Biden. In the deepfake, Biden appeared to invoke the Selective Service Act, which would initiate involuntary conscription, in order to swell American forces for military involvement in Ukraine and the Taiwan Strait [Posobiec, 2023]. What was notable about this deepfake, however, was that it didn't seem primarily intended to deceive—after Biden's speech, the video cuts to (a non-deepfaked) Posobiec himself who announces that, “What we just played for you was a sneak preview; coming attractions; a glimpse into the world beyond. Now that was an AI—I want to say re-creation, but maybe a *pre-creation*—a pre-creation of President Biden, designed and scripted by our producers here for the show, of what could happen...”. Posobiec is using synthetic media not to deceive the audience about the past, but to tell them, in vivid terms, about his vision of the future.

And a small but growing literature in computational social sciences has been experimenting with getting research subjects to use generative AI tools for envisioning the future, in both utopian and dystopian ways [e.g. Rafner et al., 2021]. For instance, in a process they called “facilitated speculation augmented with generative AI”, Epstein et al. [2022] prompted participants to write down descriptions of a re-imagined world, with results including things like, “Biophilic vertical gardens lining neighborhood roads, creating function and beautiful public spaces”, “Public spaces: solidarity-building. The intersection of oceans and relationships. Publicly accessible oceanic vistas”, and “Holistic traditional medicine as an art form”. The researchers then ran participants' responses through a text-to-image model, to create images of what the participants had imagined.

These uses of generative AI, notably, are not designed to deceive or engender skepticism: they function as vivid ways to tell others what the speaker privately envisions. While the language that participants used to prompt the models in say Epstein et al.'s work was always available for use in telling, it is clear that images sometimes have a vivacity and a power to compel or engender understanding that words don't. And while anyone with artistic abilities could always have created images based on linguistic descriptions, this ability is now available to those who lack such talent. Communicating about the future has always been possible, but vivid affirmative conceptions of how it might be have been confined, by unequal distributions of skill and by technical limitations, to non-publicity. The affordances of synthetic media around telling change this.

The communicative affordances of synthetic media are then complicated and ambivalent. What does it mean that our communicative powers should shift from the domain of showing to that of telling? Photo- and video-graphic showing has had a special relationship with the past; there have never been videos of the future, after all. And while synthetic-media-assisted telling can address past and present too, it extends special possibilities to our communications about the future. The technosocial possibility this essay has considered, then, is one in which our capacities for communication are partially redistributed from past to future. Our capacities for maintaining relational equality through

the installation of a common conception of the past have diminished. But our capacity to reach real understanding of each other's visions for the future has increased. Whether this shift in our political possibilities is for good or ill, I cannot say.

## References

- Catharine Abell. The Epistemic Value of Photographs. In Catharine Abell and Katerina Bantinaki, editors, *Philosophical perspectives on depiction*. Oxford University Press, Oxford, 2010. ISBN 978-0-19-958596-0.
- Elizabeth S. Anderson. What Is the Point of Equality? *Ethics*, 109(2):287, January 1999. ISSN 00141704. doi: 10.1086/233897.
- Daniel L Byman, Chongyang Gao, Chris Meserole, and V S Subrahmanian. Deepfakes and International Conflict. Technical report, Brookings Institute, 2023.
- Caroline Cahm. *Kropotkin and the Rise of Revolutionary Anarchism, 1872-1886*. Cambridge University Press, Cambridge, UK, 1989.
- Steve Casimiro. Can Digital Photos Be Trusted? *Popular Science*, September 2005. URL <https://www.popsci.com/scitech/article/2005-09/can-digital-photos-be-trusted/>.
- Dan Cavedon-Taylor. Photographically Based Knowledge. *Episteme*, 10(3):283–297, September 2013. ISSN 1742-3600, 1750-0117. doi: 10.1017/epi.2013.21.
- Robert Chesney and Danielle Keats Citron. Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security. *California Law Review*, 107: 1753–1820, 2018. ISSN 1556-5068. doi: 10.2139/ssrn.3213954.
- Jonathan Cohen and Aaron Meskin. On the Epistemic Value of Photographs. *Journal of Aesthetics & Art Criticism*, 62(2):197–210, 2004. ISSN 00218529. doi: 10.1111/j.1540-594X.2004.00152.x.
- Lili Corbus Bezner. *Photography and Politics in America: From the New Deal into the Cold War*. The John Hopkins University Press, Baltimore, 1999.
- Jenny L. Davis. *How Artifacts Afford: The Power and Politics of Everyday Things*. The MIT Press, August 2020. ISBN 978-0-262-35888-0. doi: 10.7551/mitpress/11967.001.0001. URL <https://direct.mit.edu/books/book/4865/How-Artifacts-AffordThe-Power-and-Politics-of>.
- Benj Edwards. AI-imager Midjourney v5 stuns with photorealistic images—and 5-fingered hands, March 2023. URL <https://arstechnica.com/information-technology/2023/03/ai-imager-midjourney-v5-stuns-with-photorealistic-images-and-5-fingered-hands/>.
- Elon Musk [@elonmusk]. @cb\_doge Yikes. Def not me., May 2022. URL <https://twitter.com/elonmusk/status/1529484675269414912>.
- Ziv Epstein, Hope Schroeder, and Dava Newman. When happy accidents spark creativity: Bringing collaborative speculation to life with generative AI, June 2022. URL <http://arxiv.org/abs/2206.00533>. arXiv:2206.00533 [cs].

- Don Fallis. The Epistemic Threat of Deepfakes. *Philosophy & Technology*, 34(4): 623–643, December 2021. ISSN 2210-5441. doi: 10.1007/s13347-020-00419-2.
- Hany Farid. Creating, Using, Misusing, and Detecting Deep Fakes. *Journal of Online Trust and Safety*, 1(4), September 2022. ISSN 2770-3142. doi: 10.54501/jots.v1i4.56. Number: 4.
- Mia Fineman. *Faking It: Manipulated Photography Before Photoshop*. Metropolitan Museum of Art, New York, 2012.
- Miranda Fricker. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press, 2007.
- Gerald Gaus. *The Order of Public Reason: A Theory of Freedom and Morality in a Diverse and Bounded World*. Cambridge University Press, Cambridge, 2011.
- James Gibson. *The Ecological Approach to Visual Perception*. Lawrence Erlbaum Assoc., 1979.
- Sanford C. Goldberg. Epistemically engineered environments. *Synthese*, 197(7): 2783–2802, July 2020. ISSN 0039-7857, 1573-0964. doi: 10.1007/s11229-017-1413-0.
- Alvin I. Goldman. Discrimination and Perceptual Knowledge. *Journal of Philosophy*, 73(20):771–791, 1976.
- Herbert Paul Grice. Meaning. *Philosophical Review*, 66(3):377–388, 1957.
- Herbert Paul Grice. Meaning Revisited. In *Studies in the Way of Words*, pages 283–303. Harvard University Press, 1989.
- James Guillaume. *L’Internationale: documents et souvenirs (1864-1878)*, volume 4. Tresse et Stocke, Paris, 1910.
- Joshua Habgood-Coote. Deepfakes and the epistemic apocalypse. *Synthese*, 201(103), 2023.
- Keith Raymond Harris. Epistemic Domination. *Thought: A Journal of Philosophy*, 2023. ISSN 2161-2234. doi: 10.5840/tht202341317.
- Patricia Hill Collins. *Black Feminist Thought: Knowledge, Consciousness and the Politics of Empowerment*. Routledge, New York, 1990.
- Robert Hopkins. Factive Pictorial Experience: What’s Special about Photographs? *Nous*, 46(4):709–731, December 2012. ISSN 00294624. doi: 10.1111/j.1468-0068.2010.00800.x.
- Nils Hulzebosch, Sarah Ibrahim, and Marcel Worring. Detecting CNN-Generated Facial Images in Real-World Scenarios, May 2020. URL <http://arxiv.org/abs/2005.05632>. arXiv:2005.05632 [cs].

- Ian Hutchby. Technologies, Texts and Affordances. *Sociology*, 35(2):441–456, May 2001. ISSN 0038-0385, 1469-8684. doi: 10.1177/S0038038501000219.
- Megan Hyska. Against Irrationalism in the Theory of Propaganda. *Journal of the American Philosophical Association*, 9:302–317, 2023a. ISSN 2053-4477, 2053-4485. doi: 10.1017/apa.2022.4. URL [https://www.cambridge.org/core/product/identifier/S2053447722000045/type/journal\\_article](https://www.cambridge.org/core/product/identifier/S2053447722000045/type/journal_article).
- Megan Hyska. Luck and the value of communication. *Synthese*, 201(3):96, March 2023b. ISSN 1573-0964. doi: 10.1007/s11229-023-04077-7. URL <https://doi.org/10.1007/s11229-023-04077-7>.
- Avery Kolers. *A Moral Theory of Solidarity*. Oxford University Press, Oxford, UK, 2016.
- Sasha Koren. Can Publishers Use Metadata to Regain the Public’s Trust in Visual Journalism?, January 2020. URL <https://open.nytimes.com/can-publishers-use-metadata-to-regain-the-publics-trust-in-visual-journalism-ee32707c5662>.
- Jennifer Lackey. The Duty to Object. *Philosophy and Phenomenological Research*, 101(1):35–60, 2020. ISSN 1933-1592. doi: 10.1111/phpr.12563.
- Federica Lago, Cecilia Pasquini, Rainer Böhme, Hélène Dumont, Valérie Goffaux, and Giulia Boato. More Real Than Real: A Study on Human Visual Perception of Synthetic Faces [Applications Corner]. *IEEE Signal Processing Magazine*, 39(1):109–116, January 2022. ISSN 1558-0792. doi: 10.1109/MSP.2021.3120982.
- Charles Larmore. *The Morals of Modernity*. Cambridge University Press, Cambridge, 1996.
- Dominic McIver Lopes. *Four Arts of Photography: An Essay in Philosophy*. John Wiley & Sons, Hoboken, 2016.
- Taylor Matthews. Deepfakes, Fake Barns, and Knowledge from Videos. *Synthese*, 201(2):41, January 2023. ISSN 1573-0964. doi: 10.1007/s11229-022-04033-x.
- Jose Medina. *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations*. Oxford University Press, 2013.
- José Medina. Resisting Racist Propaganda: Distorted Visual Communication and Epistemic Activism. *The Southern Journal of Philosophy*, 56:50–75, September 2018. ISSN 00384283. doi: 10.1111/sjp.12301.
- Raphaël Millière. Deep learning and synthetic media. *Synthese*, 200(3):231, May 2022. ISSN 1573-0964. doi: 10.1007/s11229-022-03739-2.

- Benito Mussolini. Preface. In Istituto nazionale L'unione cinematografica educativa, editor, *L'Italia fascista in cammino; 516 fotografie*. Istituto poligrafico dello Stato, Rome, 1932.
- Stephen Neale. Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15(5):509–559, 1992.
- Sophie J. Nightingale and Hany Farid. AI-synthesized faces are indistinguishable from real faces and more trustworthy. *Proceedings of the National Academy of Sciences*, 119(8), 2022. doi: 10.1073/pnas.2120481119.
- Donald A. Norman. *The Psychology of Everyday Things*. MIT Press, 1988.
- Carl Ohman. Introducing the pervert's dilemma: a contribution to the critique of Deepfake Pornography. *Ethics and Information Technology*, 22(2):133–140, June 2020. ISSN 1388-1957, 1572-8439. doi: 10.1007/s10676-019-09522-1.
- Gordon Parks. *Invisible man: Gordon Parks and Ralph Ellison in Harlem*. Steidl Books, the Gordon Parks Foundation, and the Art Institute of Chicago, Göttingen, Germany, 2016. ISBN 978-3-95829-109-6.
- Philip Pettit. *A Theory of Freedom*. Oxford University Press, Oxford, 2001.
- Francesco Pierini. Deepfakes and depiction: from evidence to communication. *Synthese*, 201(3):97, March 2023. ISSN 1573-0964. doi: 10.1007/s11229-023-04093-7.
- Jack Posobiec. Sunday Special: Ukrainistan - The Battleground of Empires, February 2023. URL <https://www.youtube.com/watch?v=htGHSJzcqWw>.
- K R Prajwal, Rudrabha Mukhopadhyay, Vinay P. Namboodiri, and C.V. Jawahar. A Lip Sync Expert Is All You Need for Speech to Lip Generation in the Wild. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 484–492, Seattle WA USA, October 2020. ACM. ISBN 978-1-4503-7988-5. doi: 10.1145/3394171.3413532.
- Janet Rafner, Steven Langsford, Arthur Hjorth, Miroslav Gajdacz, Lotte Philipsen, Sebastian Risi, Joel Simon, and Jacob Sherson. Utopian or Dystopian?: using a ML-assisted image generation game to empower the general public to envision the future. In *Creativity and Cognition*, pages 1–5, Virtual Event Italy, June 2021. ACM. ISBN 978-1-4503-8376-9. doi: 10.1145/3450741.3466815.
- John Rawls. *Justice as Fairness: A Restatement*. Harvard University Press, Cambridge, MA, 2001.
- Jorge Ribalta. Introduction. In Jorge Ribalta, editor, *The Worker Photography Movement (1926-1939): Essays and Documents*, pages 12–22. Museo Nacional Centro De Arte Reina Sofia, Madrid, 2011.

- Allissa V. Richardson. *Bearing Witness While Black: African Americans, Smartphones, and the New Protest #Journalism*. Oxford University Press, August 2020. ISBN 978-0-19-093552-8. doi: 10.1093/oso/9780190935528.003.0008.
- Regina Rini. Deepfakes and the Epistemic Backstop. *Philosophers' Imprint*, 20(24):16, 2020.
- Regina Rini and Leah Cohen. Deepfakes, Deep Harms. *Journal of Ethics and Social Philosophy*, 22(2), July 2022. ISSN 1559-3061. doi: 10.26556/jesp.v22i2.1628.
- Tom Roberts. How to do things with deepfakes. *Synthese*, 201(2):43, January 2023. ISSN 1573-0964. doi: 10.1007/s11229-023-04044-2.
- Samuel Scheffler. What Is Egalitarianism? *Philosophy & Public Affairs*, 31(1): 5–39, 2003. ISSN 0048-3915.
- Christian Schemmel. Distributive and relational equality. *Politics, Philosophy & Economics*, 11(2):123–148, May 2012. ISSN 1470-594X. doi: 10.1177/1470594X11416774.
- Tom Simonite. A Zelensky Deepfake Was Quickly Defeated. The Next One Might Not Be. *Wired*, March 2022. ISSN 1059-1028. URL <https://www.wired.com/story/zelensky-deepfake-facebook-twitter-playbook/>.
- Brian Skyrms. *Signals: Evolution, Learning, & Information*. Oxford University Press, 2010.
- Rebecca Solnit. *River of Shadows: Eadward Muybridge and the Technological Wild West*. Viking Penguin, New York, 2003.
- Dan Sperber and Deirdre Wilson. *Relevance: Communication and Cognition*. Blackwell, 1986.
- Dan Sperber and Deirdre Wilson. Beyond Speaker's Meaning. *Croatian Journal of Philosophy*, 15(44):117–149, 2015.
- P. F. Strawson. Intention and Convention in Speech Acts. *The Philosophical Review*, 73(4):439–460, 1964. ISSN 0031-8108. doi: 10.2307/2183301.
- Ruben Tolosana, Ruben Vera-Rodriguez, Julian Fierrez, Aythami Morales, and Javier Ortega-Garcia. Deepfakes and beyond: A Survey of face manipulation and fake detection. *Information Fusion*, 64:131–148, December 2020. ISSN 15662535. doi: 10.1016/j.inffus.2020.06.014.
- Zeynep Tufekci. *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. Yale University Press, New Haven and London, 2017.
- Olúfémí O. Táíwò. *Elite Capture: How the Powerful Took Over Identity Politics (And Everything Else)*. Haymarket Books, 2022.



- Scott Walden. *Truth in Photography*. Blackwell Publishing Ltd, Oxford, UK, 2008. ISBN 978-0-470-69665-1 978-1-4051-3924-3. doi: 10.1002/9780470696651.ch4.
- Kendall L. Walton. Transparent Pictures: On the Nature of Photographic Realism. *Critical Inquiry*, 11(2):246–277, 1984. ISSN 0093-1896.
- Lani Watson. *The Right to Know: Epistemic Rights and Why We Need Them*. Routledge, 2022.
- Jonathan Wolff. Fairness, Respect, and the Egalitarian Ethos. *Philosophy & Public Affairs*, 27(2):97–122, 1998. ISSN 0048-3915.
- Amy Louise Wood. *Lynching and Spectacle: Witnessing Racial Violence in America, 1890-1940*. University of North Carolina Press, Chapel Hill, 2009. ISBN 978-0-8078-7811-8.
- Iris Marion Young. *Justice and the politics of difference*. Princeton University Press, Princeton, N.J, 1990. ISBN 978-0-691-07832-8.
- Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-Shot Adversarial Learning of Realistic Neural Talking Head Models. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9458–9467, October 2019. doi: 10.1109/ICCV.2019.00955. ISSN: 2380-7504.
- Ruihan Zhang. MIT AI for Filmmaking Hackathon 2023 Brings Dreams to Life, February 2023. URL <https://www.media.mit.edu/posts/mit-ai-for-filmmaking-hackathon-2023-brings-dreams-to-life/>.