

DEEPAKES, PORNOGRAPHY AND CONSENT

Abstract: Political deepfakes have prompted outcry about the diminishing trustworthiness of visual depictions, and the epistemic and political threat this poses. Yet this new technique is being used overwhelmingly to create pornography, raising the question of what, if anything, is wrong with the creation of deepfake pornography. Traditional objections focusing on the sexual abuse of those depicted fail to apply to deepfakes. Other objections—that the use and consumption of pornography harms the viewer or other (non-depicted) individuals—fail to explain the objection that a *depicted* person might have to the creation of deepfake pornography that utilises images of them. My argument offers such an explanation. It begins by noting that the creation of sexual images requires an act of consent, separate from any consent needed for the acts depicted. Once we have separated these out, we can see that a demand for consent can arise when a sexual image is of us, even when no sexual activity was actually engaged in, as in the case of deepfake pornography. I then demonstrate that there are two ways in which an image can be ‘of us’, both of which can exist in the case of deepfakes. Thus, I argue: if a person, their likeness, or their photograph is used to create pornography, their consent is required. Whenever the person depicted does not consent (or in the case of children, can’t consent), that person is wronged by the creation of deepfake pornography and has a claim against its production.

Keywords: consent, deepfakes, doctored, pornography, privacy, photography

1 Introduction

The media is awash with panic about deepfakes: images, still or moving, created with the help of artificial intelligence (AI) that take a photograph or footage of an existing person and modify it to make it appear as if they did or said something they didn't. Fake footage of politicians and cultural figures now abound: from Barack Obama and Mark Zuckerberg to Jon Snow's moving apology for Season 8 of *Game of Thrones*. Because of their use, especially in the political sphere, discussions have focused on the threat deepfakes poses to trust, social cohesion, and democracy (Fallis 2020; Chesney and Citron 2019; Floridi 2018; Rini 2020; Carlson 2021). However, Jeffrey Shallit's 'First Law of New Media' states that every new medium of expression will be used for sex (Shallit 1996). Deepfake technology is no exception.

In fact, deepfakes shot into the public awareness because of their use in creating pornography. In 2017, the online magazine *Motherboard* broke a story about a Reddit user called 'deepfakes' who had uploaded doctored pornography of Scarlett Johansson, Taylor Swift, and Gal Gadot (amongst others).¹ Their likenesses had been superimposed onto porn performers using AI to render the final pornographic videos realistic at first glance. The cybersecurity company Deeptrace found that 96% of deepfakes were pornographic and that 99% of those mapped women's faces onto those of porn actors (Ajder et al. 2019).

While there has been public outcry about the creation of deepfake pornography, there has been far less sustained attempt to explain exactly why their creation is objectionable², in particular from the point of view of those whose images are used to

¹ <https://www.vice.com/en/article/gydydm/gal-gadot-fake-ai-porn>

² Exceptions include a 2021 paper by de Ruyter (2021) and a publicly available forthcoming paper by Rini and Cohen (2022).

create them. This paper addresses this gap, providing an argument that explains the objection that individuals have against the production of deepfake pornography when images of them are used.

The structure of this paper is as follows. I begin, in §2, by giving an account of deepfake pornography and raising the central question of the paper: on what grounds can those whose images are used to create deepfake pornography object to its creation? In §3, using the case of deepfake child pornography as a case study, I uncover and reject the assumption that makes this question seem puzzling: that the only way the creation of pornography wrongs the depicted person is because it involves sexual abuse in its creation.

In §4, I establish that there are two ways in which an image can be 'of us' and that these two ways can apply even when the final image is a deepfake. I conclude, in §5, by presenting my consent-based objection to deepfake pornography, answering the question I started with. I explore two objections that might be raised and discuss how my argument can be extended to provide a basis for objections to the creation and use of other kinds of images in contexts beyond deepfake pornography.

2 Existing Arguments against Deepfake Pornography

2.1 Defining Deepfake Pornography

To begin to answer the question of how deepfake pornography might wrong the depicted person, we first need to know what it is. Part of the complexity of giving a definition of deepfake pornography lies in the fact that it is notoriously difficult to give

a definition of pornography.³ For the purposes of this paper, I assume that we have a handle on what counts as *photographic* pornography.⁴ Deepfake pornography, then, utilises photographs (a term I use to include photographic footage) of actual persons (including non-sexual photographs) via deepfake technology to create a final (false or misleading) image that, were it a photograph, would be classed as pornography.⁵

2.2 Existing Arguments

Despite the fact that, in deepfake pornography, the person depicted as engaging in sexual activity did not in fact do as they are depicted to have done, they are often nonetheless indistinguishable from *photographic* pornography. As such, answers to the question of what is wrong with deepfake pornography have tended to focus on the harms arising from its use and consumption.

Deepfake pornography can clearly be *used* in various ways that are morally wrong, for example grooming a child for abuse (Adelman 1996; Armagh 2002; Bergelt 2003). However, in these cases, the wrong of the *use* is nothing to do with the image itself: it is straightforwardly wrong to groom a child, no matter what is used to do so, be it

³ There are a variety of ways of defining pornography (Williams 1981; Dworkin 1985; Longino 1980; Mackinnon 1987; Rea 2001). Defining child pornography is even more complex especially in the digital age (Taylor and Quayle 2003; Benn 2019).

⁴ Note that there is an important difference between images that are 'photographic' and those that are 'photographs' as well as between those that are 'pornographic' and those that are 'pornography' (Patridge 2013). However, for brevity, I treat these terms as equivalent. Thus, by 'photographic pornography' I mean pornography that are also photographs, and by 'pornographic photographs', I mean photographs that are also pornography.

⁵ The term 'images' here is intended to refer to both still and moving images. Note that there is a wide range of terminology that has been used in this space, especially in the legal tradition. What I am calling doctored images have been called 'pseudo-photographs' (Strikwerda 2011, 140) or 'morphed' images (Ashcroft vs. The Free Speech Coalition 2002, mentioned in the Opinion of the Court delivered by Justice Kennedy; Levy 2002, 319; Krone 2004; Eneman, Gillespie, and Stahl 2009; Burke 1997, 440; Bergelt 2003, 570; Armagh 2002, 1994). I avoid the term 'pseudo-photographs' because it has also been used to refer to any image (whether doctored or completely CGI'd) that is indistinguishable from a photograph (see for example the UK law concerning 'Indecent Images of a Child'). I avoid the term 'morphed' as it is best reserved for one specific way in which images can be doctored (Farid 2004).

deepfake child pornography, photographic child pornography, adult pornography, non-sexual images, or objects such as sweets or toys.

Let us limit our discussion to the private and personal use of pornography, for example, as a means of achieving sexual arousal. In this case, three arguments can be raised against the *consumption* of deepfake pornography. The first is to extend the argument that viewing pornography harms the *consumer*, causing them to be deprived or corrupted. This is at the heart of the obscenity objections to pornography that have dominated legal discussions (Koppelman 2005; Henkin 1963; Regina v. Hicklin, L. R. 3 Q. B. 1868). The second is that, like photographic pornography, deepfakes can encourage harm to *others*. As some have argued, violent adult pornography encourages viewers to act in ways or express views that tolerate or promote violence against women (Eaton 2007; Longino 1980; Mackinnon 1987). This argument does not rely on any specific connection between the final image and actual persons depicted: it is the resulting image's reception and role in normalising and inciting sexual violence that is the cause of the harm, and this applies as much to deepfake pornography as it does to photographic pornography. This argument—that pornography can lead to greater tolerance or even promote sexual violence—has also been raised with respect to child pornography where many have argued that deepfake child pornography encourages paedophiles to sexually abuse non-depicted children (Taylor and Quayle 2003; Strikwerda 2011). The third argument is that deepfake pornography, just like photographic pornography, harms women as a group. For example, Carl Öhman argues that “The consumption of Deepfakes is undeniably a highly gendered phenomenon, and arguably plays a role in the social degradation of women in society” (Öhman 2020, 139). This draws on the argument that pornography can objectify or subordinate women by sexualising their inequality (May and Friedman 1985; Langton 1993; Mackinnon

1987).⁶ This argument about the sexualisation of inequality has also been applied to synthetic *child* pornography (Levy 2002).⁷

Deepfake pornography may well be wrong in the ways these arguments suggest. In fact, I am deeply sympathetic to many of these arguments and in no way seek to reject them. However, regardless of the plausibility of these existing arguments, they fail to explain how deepfake pornography can wrong *those depicted*. My argument addresses this question directly and therefore differs from existing arguments in three important respects. First, existing arguments locate the wrongs of deepfake pornography in the *use* or *consumption* of such images. I establish what is wrong with *producing* deepfake pornography. Second, existing arguments are based on empirical claims concerning the consumption of pornography and harm. My argument does not stand or fall depending on the outcome of empirical research, which is particularly important because the causal connection between the consumption of pornography and harm has been notoriously difficult to definitively establish. And finally, existing arguments focus on harms to non-depicted persons (either as individuals or as a group). My argument brings to light how deepfake pornography wrongs the *depicted* person.

⁶ Note that these positions are bolstered when feminists restrict their argument to inegalitarian adult pornography (Eaton 2007).

⁷ Levy argues against the production of virtual child pornography because it sexualises inequality and this harms *women*. Thus, it has been critiqued as only explaining the wrong of virtual child pornography *indirectly*: he doesn't explain how it wrongs *children* (Patridge 2013).

3 *Photographic Pornography and Consent*

3.1 The Puzzle of Deepfake Pornography

So: what, if anything, is wrong with deepfake pornography such that the depicted person has a specific grievance against its production? Let's start by considering how the creation of *photographic* pornography wrongs the person depicted.

A key argument is that the creation of photographic pornography wrongs the depicted person when it harms that person. While this argument has been made by prominent theorists about the harms done in the adult porn industry (MacKinnon and Dworkin 1997; Andrea Dworkin 1981; Lovelace and McGrady 1980), it has dominated discussions of child pornography. In *New York v. Ferber*, for example, part of the rationale for extending prohibitions of child pornography was that it is “intrinsically related to the sexual abuse of children”, as the production of the material “requires the sexual exploitation of children” (New York v. Ferber 1982, vol. 81–55, secs. 458 U.S. 747, 760). In more recent years, many organisations and academics have argued that, instead of ‘child pornography’, we should call such images ‘child abuse images’ ([INTERPOL](#)) or ‘images of sexual abuse’ (Taylor and Quayle 2003, 7), or ‘child sexual abuse material’ ([Technology Coalition](#)). The proponents of this terminological change argue that these terms capture the real wrong of these images and express unambiguously “the nature of child pornography” (Taylor and Quayle 2003, 7). A child engaging in a sexual act is a child who is being sexually abused and a photograph of this act is simply the recording of an abusive act. Thus, these photographs are, as the National Association for People Abused in Childhood puts it, “crime scenes” (NAPAC 2016, 4). In this sense, the *photograph* per se is immaterial: the moral wrong is in the

act depicted and the fact that this wrong had to take place for the photograph to be made is what makes the production of the photograph morally wrong.⁸

And this is the puzzle of deepfake pornography: the immediate reasons we have to object to the creation of photographic pornography simply fail to apply in the case of deepfake pornography. In the case of children, sexual abuse of the depicted child is not a wrong that can be attributed to the creation of deepfake child pornography, by its very definition. If it is the *only* way the creation of photographic child pornography wrongs depicted children, then the project of finding a wrong that *deepfake* child pornography commits against depicted children is a non-starter. However, this position relies on the assumption that all child pornography depicts acts of child sexual abuse, an assumption that is false even if we limit our discussion to *photographic* child pornography. Consider, for example, a photograph of a child masturbating or of their genitals. This photograph does not depict an *act* that constitute child sexual abuse and yet would intuitively and, in many jurisdictions, legally be classed as child pornography.⁹

Thus, the production of photographic child pornography does not necessarily involve child sexual abuse in the acts depicted. The assumption that it did was the main reason to be sceptical of the claim that deepfake child pornography could wrong the children depicted in it. Without it, two questions remain: what is wrong with all photographic child pornography if not child sexual abuse? And does this wrong apply to deepfake child pornography? Answering these two questions will help us ascertain

⁸ This argument has also dominated anti-porn feminist critiques that have focused on the harm done to women actors in the production of pornography, as documented in Lovelace and McGrady (1980).

⁹ It would also legally be classified as child pornography in many jurisdictions despite not depicting child sexual abuse (for the US, see (New York v. Ferber 1982; Burke 1997, 442, footnote 13); for the UK, see the SAP and COPINE scales (Taylor, Holland, and Quayle 2001, 101).

how deepfake images more generally can wrong the persons (children and adults) whose image is used. Let's turn to answering the first of these questions.

3.2 Consent, Sex Acts and Sexual Images

The examples above demonstrate that the ethical status of an act is distinct from the ethical status of producing an image of that act. Thus, the permissibility of an act does not entail the permissibility of making an image of that act. For example, it is permissible for someone to give witness testimony; however, it is in general impermissible to take a photograph of a witness in a courtroom (it is illegal to do so in the UK under The Criminal Justice Act 1925 Sec. 41). Conversely, the *impermissibility* of an act does not entail that producing an image of that act is also impermissible. For example, there are many occasions on which it is permissible to photograph morally impermissible acts, for example of war crimes, police violence or domestic abuse. So, what makes the production of a sexual image permissible or impermissible, if it isn't determined by the act depicted? My answer is *consent*.

Consent has long been a cornerstone of explanations of what makes a sexual act permissible in one context but impermissible in another. However, while discussions of consent have dominated the discussion of sexual *activity*, it has been less the focus of discussions about the creation of sexual *images*. Just as a sexual act can be permissible if all parties consent but is impermissible if they do not, it can also be permissible to make an image of people engaging in that sexual act if they consent, but it is impermissible if they do not. Importantly, the above discussion brings home the fact that consenting to engaging in a sexual act does not entail consent to the production

of an image of that act. A separate act of consent is required.¹⁰ Consider cases where pornography is made of someone without their consent: for example, where spy cameras have captured and distributed footage of people having sex in Airbnbs. The impermissibility of this lies not in any impermissibility of the sexual activity but in the lack of consent to the *recording* of that sexual activity.

Thus, just as sexual activity creates a special demand for consent, so too does the creation of sexual photographs. As children cannot consent, photographic pornography *necessarily* wrongs the children depicted.¹¹ With respect to adults, I will assume that there are circumstances where adults can consent to the production of photographic pornography of themselves but that when an adult does not consent, the creation of photographic pornography wrongs them.

However, this does not yet explicitly speak to the wrong, if any, of *deepfake* pornography. In the next section, I discuss the two ways in which we are connected to images of us—connections that ground the demand for consent—and show that these connections persist even if an image is doctored, as in the case of deepfakes.

¹⁰ Note that, nevertheless, consent might not be *sufficient* to make the production of pornography permissible as the other arguments (discussed in §1) against the production and consumption of pornography—that it harms non-depicted individuals or groups or sexualises inequality itself—apply and render the creation of all pornography (or certain types of pornography) impermissible *even if* those depicted consent.

¹¹ For a discussion of the capacities needed for sexual consent, see (Archard 1998). Note that I leave aside whether the capacities needed to have the ability to consent in the case of sexual activity and the production of sexual images arise at the same age, something rejected by most Western countries, which legally define a child as under 16 for the former but under 18 for the latter (Healy 1997). Some might object, assuming that if someone has the capacities necessary to permissibly engage in a sexual act, they must also have the capacities necessary to consent to the production of an image of that act. However, this overlooks the fact that there are sexual acts, such as masturbation, that have no age of consent. Infant masturbation has been recorded as early as 2 months old (Hansen and Balslev 2009) and yet it cannot be the case that a 2-month-old has the capacities necessary to be able to consent to anything at all. Also note law-specific definitions of what it is to be a child are common in areas of activity unrelated to sex (for example, voting and criminal responsibility) (“The UN Convention of the Rights of the Child” 2010, Article 1).

4 *Images of Us*

4.1 Identifiability

So: what is it about photographs that give them a special connection to the depicted? One obvious answer is that photographs tend to look like us. Insofar as we are identifiable from photographs, and these images say something about us, we have reason to be morally concerned about them.

The intuitive concern about identifiability grounds the current U.S. legal position on child pornography. In 1996, it was already illegal to create, possess, and distribute any visual depiction that involved an actual child engaging in sexually explicit conduct. The definition of child pornography was later expanded to include any visual depiction (including created or doctored ones) in which it appears that an *identifiable* child is engaging in sexually explicit conduct.¹²

One question that arises immediately is: identifiable to whom? The U.S. legal position is that the child must be “recognizable as an actual person by the person’s face, likeness, or other distinguishing characteristic, such as a unique birthmark or other recognizable feature” (*Child Pornography Prevention Act of 1996 (CPPA)* 1996, sec. 18 U.S.C. §2252A). This definition is both very narrow and extremely subjective. Rarely are birthmarks so unique (or known to be) as to identify someone. Also, there are certain things from which only people who know us well may be able to recognise us.

What is important to note is that, no matter how we define it, the property of identifiability is not limited to photographs: many deepfake images maintain the

¹² This was the outcome of the Child Pornography Prevention Act (*Child Pornography Prevention Act of 1996 (CPPA)* 1996, Pub. L. No. 104-208, codified at 18 U.S.C. §2252(8)C.) and its subsequent challenge in (*Ashcroft vs. The Free Speech Coalition* 2002).

identifiability of the depicted. If *photographic* pornography is objectionable in part because it depicts identifiable persons, then there is reason to object to *any* deepfake pornography that depicts identifiable persons. Thus, whatever explains the problem with identifiability should explain objections to both photographic and doctored pornography that depict identifiable persons. I turn now to arguing that the best way to understand what is objectionable about the use of someone's identifiable image is that it involves a violation of consent.

4.2 Defamation, Privacy and Consent

Philip Brey outlines two reasons to care about identifiable images of us: defamation and privacy (Brey 2008).¹³ These may well provide the best basis for a legal restriction on the depiction of real people in certain cases; however, using deepfake child pornography as the test case, I show that neither defamation nor privacy completely capture what is *morally* wrong with doctored pornography that depicts identifiable persons.

Let's begin with defamation. There is clearly something behind the concern that deepfake images that depict identifiable children tell a lie about those children. However, defamation requires not only that lies are told but that the public reputation of that person is damaged because they are depicted as *morally depraved or ridiculous* (Brey 2008, 11–12). This cannot be our objection to child pornography. An image of a child masturbating does not depict that child as morally depraved or ridiculous. Neither does an image that depicts child sexual abuse. It would defame an adult to falsely

¹³ I leave aside the third reason he discusses: publicity. The right of publicity is usually reserved for celebrities who make their living partly from this commercial use. My argument establishes the interests of *all* of us in not having our images used in the creation of pornography without our consent, even if we are not celebrities.

portray them as the *perpetrator* of child sexual abuse; but any temptation to see such images as defaming *the child* simply reveals something deeply concerning about our current culture around sex, shame, and victim-blaming.¹⁴

Concerns about privacy seem to get closer to the mark. Pornography often depicts something private: a private act or a private body part. This can explain what is wrong with photographic child pornography that depicts acts that are not themselves constitutive of child sexual abuse. A photograph of a child masturbating or naked, it can be argued, violates that child's privacy by intruding on their private affairs. However, as noted before, whatever explains the problem with identifiability should be able to explain our objection to both photographic and deepfake pornography that depicts identifiable persons. And the argument from privacy cannot explain what is wrong with deepfake pornography where a person is identifiable in the final image: the *original* photograph has no sexual content and so does not constitute a violation of privacy (let's stipulate that the images used are publicly available). The final image has sexual content and so *looks* like a violation of privacy, but no such violation has in fact taken place: the violation is merely simulated. Deepfake pornography that appeared to be of you performing sexual acts might certainly *feel* like a violation. But if there is a violation, it is not one of privacy as your *actual* private affairs are not intruded on at all.

¹⁴ But perhaps 'ridiculous' here should be understood not as a normative term—deserving of ridicule—but as descriptive—likely to be ridiculed. And indeed, many things unworthy of ridicule are likely to be ridiculed, such as an image of a celebrity on the toilet. However, even if we were cynical enough to believe that people *would* ridicule a child for being depicted as the victim of child abuse, it would be strange if our objection to a child being identifiable in pornography is grounded in the fact that they will be ridiculed as a matter of fact, as then the objection would disappear if people rightly came to see them as the victims they are. And surely our objection is not as fleeting as that.

Another possibility is that identifiable deepfake pornography exposes those depicted to a risk of stigma. For example, deepfakes could be made of politicians appearing to engage in sex with men in counties where attitudes towards homosexuality are negative and intolerant. Despite the fact that gay sex is, in fact, not depraved, this would constitute a serious reputation harm to those depicted. Nevertheless, this argument from the harm or risk of stigma is not sufficient because it does not explain why an identifiable deepfake of you can be problematic even if it only depicts you having the kind of sex that is in no way stigmatised in your current society.

Furthermore, the concern with privacy is itself, at its base, an issue of consent. We often allow people into our intimate space such that we enjoy *less* privacy. What is of moral concern is not a loss of privacy but a *violation* of privacy (Introna 1997, 262). And loss of our privacy is a violation of our privacy when it occurs *without our consent*.

My contention is that what is wrong with the production of pornographic photographs in which someone is identifiable is that the production of such images requires consent and, when that person has not consented (or in the case of children, *cannot* consent), they are wronged by the production of such photographs. As it is the *identifiability* of an individual in an image that grounds the requirement for consent in the case of photographs, it follows that *deepfake* pornography depicting an identifiable person also requires consent. Thus, when an adult does not consent to their likeness being used in the creation of deepfake pornography, the production of such doctored images wrongs them. And in the case of children, just as they cannot consent to the creation of *photographic* pornography of themselves, they also cannot consent to having their likeness used to create pornography. Thus, whether or not the final image is a deepfake, whenever pornography is made that depicts an identifiable child, that child is necessarily wronged.

4.3 Two Connections

I could stop here. Almost all deepfake pornography depicts an identifiable person. The entire point of using the images of celebrities is their recognisability. That those depicted are identifiable is part of the intended harm to victims in the case of 'revenge

porn'.¹⁵ However, I want to also address the claim of someone whose starting image is used even if they are rendered *unidentifiable* in the final image.

To establish my consent-based objection to deepfake pornography even when the person is not identifiable in the final image, let us return to photographs and our connection to them. While it is true that we are often identifiable in photographs, there is another connection that we have to photographs of us.

To understand this second connection, let's start by exploring what makes photographs special. The idea that photographs bear a special relationship to what they depict is noted in the well-documented sense of nearness, intimacy, contact or proximity that we have when looking at photographs. In the words of Elizabeth Barrett Browning, writing only a few years after the invention of photography: it is "not merely the likeness which is precious in such cases—but the association and the sense of nearness involved in the thing." (quoted in Sontag 2005, 143). Robert Hopkins explains this "distinctive power" of photographs in terms of putting us "in a relation to their objects that is somehow more intimate, more direct, than that in which we stand to the objects handmade pictures depict" (Hopkins 2012, 709).

So, what underpins this special relationship of nearness or proximity? Drawing on Mary Ann Doane's terminology, I call it a 'material connection' (Doane 2007). Doane's work is based on C.S. Peirce's semiology, according to which, the relationship photographs have to what they depict is not only iconic (a connection based on resemblance and identifiability) but also indexical (a connection based on this material connection) (Peirce 1931, sec. 2.281).

¹⁵ While this term is in common use to capture those pornographic images taken (or shared) without the consent of one or more of those depicted, often by someone known to the person, it should be noted that not all image-based abuse of this kind has to be motivated by 'revenge' as the term 'revenge porn' suggests.

The idea that photographs or recordings in general have a material connection to what they represent “is as old as recordings themselves” (Carlson 2021, 3). This material connection underpins many prominent accounts of photography (e.g. Scruton 1981), including Dawn Phillips’ where a photograph is a record of a ‘photographic event’, that is the recording of a light image (Phillips 2009). It is also a key part of some of the most influential accounts of the phenomenology or subjective experience of photography, such as Kendall Walton’s ‘transparency thesis’ (Walton 1984), as well as the Trace Theory (Pettersson 2011; Currie 1999).¹⁶ As many have argued, it is not the epistemic status of the *content* depicted that makes photographs special (Pettersson 2011). Photographs of loved ones are often valued, even if little is learned from them (Walton 1984, 253). It is the connection to the person depicted that we value (Phillips 2009; Carlson 2021; Benovsky 2016). And this connection does not require identifiability. As Susan Sontag has argued even if a photograph of Shakespeare were “faded, barely legible, a brownish shadow, we would probably still prefer it to another glorious Holbein. Having a photograph of Shakespeare would be like having a nail from the True Cross.” (Sontag 2005, 120)

Identifiability seems like the more obvious—and therefore more important—connection that we have to photographs. Thus, it might seem less plausible that this material connection could ground a requirement for consent. However, this material connection is key.

To see why, first note that identifiability without *any* kind of connection may not be sufficient to ground a requirement for consent. Take Hilary Putnam’s classic

¹⁶ A ‘trace’ is one form of ‘index’. For example, there are marks that someone has made: footprint or a death mask are traces of, respectively, someone’s foot and someone’s face (Sontag 2005, 145). There are also transformations, for example, a burnt tree stump is a trace of a fire (Kausser 2007, 59). Note that material or indexical connections can take many forms: it could be causal (as in the case of smoke being a sign of a fire) or it could be non-causal (such as the Pole Star as an index for the North Pole) (Gouge 1965, 54–56).

example of an ant crawling around on a patch of sand, where the lines it happens to trace ends up looking like a “recognisable caricature of Winston Churchill” (Putnam 1979, 1). Putnam asks whether the ant has traced a picture that *depicts* Churchill and answers that it has not, because the ant has not seen Churchill (or a picture of Churchill) and has no intention of depicting him. Thus, although the image is identifiable as Churchill, it does not depict him because it lacks the appropriate causal connection. Identifiability matters in the case of deepfake pornography because these images are derived from photographs that do have a material connection to the depicted.¹⁷ What is significant is that it is only in *conjunction* with this other connection that our identifiability in photographs becomes important.¹⁸

Second, not only is mere identifiability insufficient, it is also not *necessary* for a connection between the depicted and the depiction, because the material connection we have to photographs can ground a requirement for consent even when we are not identifiable. Suppose someone takes a photograph of you from the neck down, or of your feet, or your genitalia. Or suppose the photograph is blurry or overexposed or only shows you in silhouette. Even you may not be able to identify yourself from these photographs, let alone a stranger. Nevertheless, they would still be photographs of you.¹⁹

¹⁷ These connections explain why we might be pulled in two directions when it comes to the interesting (albeit unlikely) case of one identical twin being involved in producing pornography (this is part of an episode of *Friends*, but I have yet to hear of a real-life case). We might understand the other twin’s objection, which is based on identifiability: that someone might believe that it was them. However, we might also understand the original twin’s argument that the objecting twin has no grounds to object because they lack the proper material connection, as they were not in fact photographed.

¹⁸ We might think that an intentional connection is enough (rather than the material one that applies in the use of photographs). This would ground objections to cases where someone is depicted intentionally in completely created synthetic pornography, which did not use a photograph of the person it depicts. I do not settle this specific case here (though it should be noted how currently and increasingly rare and infrequent such completely synthetic images are, given the availability of deepfake software).

¹⁹ Benovsky suggests that *typically* photographs fail to “depict reality as it is (they only depict things from one side, they can involve distortions, blurred background, etc.)” (Benovsky 2016, 77–78). Phillips also says “‘Photograph of’ picks out a causal relation to the objects and sources that were causally

Some might push back at the idea that a material connection without identifiability of any kind grounds a requirement for consent. Here are two counterexamples.²⁰

The first is that the *photographer* is just as causally important to the existence of the final image as the person depicted. However, it would be unusual to think that the photograph was equally a photograph *of* the photographer as of the photographed.

The second is that, outside of photographs, we might have causal connections to things that it might seem odd to think that would require my consent to be used in pornography. For example, imagine that you have sex in an alley behind my house and leave a swirled handprint in the drying paint on my fence. If I use this handprint in a pornographic image, does the causal connection that exists between you and that image, via the handprint, mean that your consent is needed?

These are thorny issues. I do not think that there is a clear answer with which all would agree. I would say that, when the material connection exists, the more that an image resembles us (with identifiability in the strong, narrow, legal sense at one end of the spectrum), the stronger claim we have against its use to create pornography. However, the fact that an image does not render us identifiable does not mean that we have no claim. Our depiction matters, even if the final image is (as noted above) blurry, grainy, or of parts of our bodies from which no one could identify us. The key sense of nearness or proximity that viewers have to those depicted in an image remains even when the image does not entail high-fidelity visual resemblance. This is sufficient for my argument.

responsible for the light image. Being a photograph of these things does not entail visual resemblance.” (Phillips 2009, 339)

²⁰ My thanks to the reviewers for bringing these interesting cases to my attention.

Moreover, it is not implausible that the persons in the cases above (the photographer or the leaver of the swirled handprint) *may well understandably* object to their causal creations being used for certain purposes. Imagine you find out that your artistic Instagram photographs of tidying hacks for children's playrooms had been used extensively as backgrounds for deepfake child pornography. Or imagine that the handprint you left in the paint ended up being used as the symbol for a racist terrorist organisation who thought that the smeared white handprint aesthetically captured their cause. You might indeed be horrified and vehemently object to these causal traces of yourself being used in these ways.

For now, however, I will set aside whether purely causal connections that are not manifested in what is captured *in* a photograph can ground requirements for consent. Instead, let's focus on the fact that this material connection exists for all photographs and that it is this material connection to those photographs that explains why the production of pornographic photographs wrongs the person depicted, even if they are not identifiable in that photograph. But does this material connection that exists between a person and a photograph of them persist through doctoring as in the case of deepfakes? I now show that it does.

4.4 The Line between Photographs and Deepfakes

To accept that photographs have the right kind of connection (established above) but deny that doctored images like deepfakes do, a clear separation must exist between them. But photographs have always been doctored. Many early photographs were a product of 'combination printing' (originally proposed by Hippolyte Bayard, as early as the 1850s) where two or more negatives were used to create a single image, in part because the lack of light sensitivity meant that, by the time the main subject of a

photograph was properly exposed, other elements, such as the sky, were bound to be over-exposed. Throughout the Victorian era the doctoring of run-of-the-mill portraits at commercial photography studios was commonplace, including the removal of wrinkles and freckles as well as the slimming of waists. It is described and recommended in several guides and instruction manuals from that era (Johnson 1898; Schriever 1909).

The digital age has put further pressure on maintaining a distinction between a photograph and a doctored image. Digital images are just data and this data is highly malleable. A digital photograph of you displayed on a low-resolution screen is just as much a photograph of you as the one displayed on a high-resolution screen, even though they are qualitatively different.²¹ Can one really be said to be the original and the other doctored?

There are plenty of ways of doctoring photographs: obscuring or exaggerating image details, altering the colour saturation and contrast, the brightness, sharpening, cropping, making it hazy, blurry, in sepia, grayscale, or pixilated, or even make it look hand-drawn or like a cartoon (Farid 2004). We can now make any of these changes at the touch of a button, just as easily as taking a photograph, through graphic filters readily available on smartphones. In fact, these graphic filters can now be applied *pre*-production: the image is previewed and taken with the filter already applied (via apps such as Facetune, Reface and MSQRD). In such cases, there is no 'original' photograph.

It is therefore not possible to maintain a sharp enough distinction between a photograph and a doctored image to ground a requirement for consent in the case of the former but not in the case of the latter: if a graphic filter is added to a pornographic

²¹ Elsewhere, I discuss how digital photography demands a conceptual change to our understanding of an image (Benn 2019).

photograph (or added before the image is even taken), the requirement for consent does not vanish.

Of course, there may come a point where, because the changes made are so many or so radical, this chain between the final image and the person in the original photograph is broken.²² Suppose just one pixel from a photograph of you is used as the basis of the skin colour of an image of someone having sex. The link in this case is perhaps not strong enough to merit a requirement for your consent. I leave the question of exactly where the line should be drawn for future research.²³ The idea that the connection we have to a final image might breakdown at some point does not undermine the more general point that the link between you and an image can remain intact through doctoring and that this connection can ground a requirement for consent when your image is put to use as pornography.

4.5 Consent and the Creation of Pornography

Doctoring a photograph cannot *remove* any original demand for consent, whether or not we are identifiable in the final image. However, it might be questioned whether *introducing* sexual content or context through doctoring an otherwise non-sexual image cannot give rise to a requirement for consent. I argue that it can. To see why, we must return to the question: what is pornography?

²² This chain-line connection between us, photographs, and doctored images of us is discussed further in (Mcmullan 2011; Poremba 2011).

²³ A place in which this distinction is likely to be key is in the emerging generative AIs such as DALL-E, Sora and Gemini which create synthetic media but by using existing images, including photographs of existing people. My argument can explain the objection that those of whom *identifiable* pornographic images are made through such means have. If we assume the chain is never broken, my argument would potentially ground the objections of anyone whose images are used to train these models when they are used for the generation of *any* pornography.

Michael Rea identifies a curious phenomenon: a naked photograph of Marilyn Monroe would be pornography if it appeared in *Hustler*, but it wasn't when it appeared in *Life* (he noted that it *was* considered pornographic when it originally appeared in a calendar in 1940, being banned in two states) (Rea 2001, 118). He asks how this can be accounted for: how an image could be pornographic even though there is no change in content. His argument, and his subsequent definition of pornography, is that some material, x , is pornography when "it is reasonable to believe that x will be used (or treated) as pornography by most of the audience for which it was produced" (Rea 2001, 120).²⁴

The advantage of this position is that, because specific kinds of *content* (such as nudity or depictions of sexual acts) are not necessary for an image to be pornography, it can account for a wide range of intuitively pornographic images, such as those depicting feet and shoes when created for shoe-fetishes (Rea 2001, 122).²⁵ This is in keeping with my argument that not all child pornography falls under the description of 'child abuse images' because this latter definition is too narrow in terms of what content an image must have in order to be child pornography. More significantly, it can explain how the *introduction* of content and alterations in context can change the audience for whom it is intended, and the reasonable expectations of how that image will be treated, and thus whether or not it is pornography. This includes its posting in certain spaces, alongside other pornographic content.²⁶ But it, arguably, also includes deepfaking

²⁴ He offers a complex definition of what it means to 'treat something as pornography' but a common-sense understanding will do for the purposes of this paper.

²⁵ It can also account for why some images with certain content, in particular children in the nude, might *not* be considered child pornography in certain contexts. For example, Spencer Elden, the naked baby on the cover of the album *Nevermind*, sued Nirvana, claiming it was child pornography. An application of Rea's argument would deny this claim, as it is unreasonable to believe that the album cover would be used or treated as pornography by most of the Nirvana fans for whom the album was produced.

²⁶ Elsewhere I offer an account of how metadata could be included as part of the context that determines if an image is pornography or not (Benn 2019).

someone's image onto footage that is pornographic. Adding sexually explicit material to someone's image *recontextualises* that individual's image such that it is reasonable to expect it will now be treated as pornography. As such, it makes a non-sexual image into pornography, in the same way that placing Marilyn Monroe's picture in *Hustler* (or a non-sexual image of a child alongside sexually explicit images of that child²⁷) turns it into pornography.

My argument goes beyond Rea's both in considering the case of doctored and deepfaked images as well as demanding that we look at how this recontextualising affects the demand for consent. My argument is that, just as our consent is needed for *content* that makes it reasonable to believe that the resulting image will be treated as pornography, our consent is also needed when our image is recontextualised such that it *becomes* pornography. This includes a picture of someone in a bikini or at a nudist beach (neither of which are necessarily pornographic) that is placed on a website containing hard-core pornography. It also addresses cases where elements are added to an image of someone such that it is likely to be treated as pornography.²⁸ This is exactly what deepfake pornography does. Whenever pornography is made of us, our consent is needed. Deepfakes are no exception.

²⁷ Taylor and Quayle discuss the real case of this sort: a 14-year-old child who, in 2000, had been depicted in large amounts of child pornography as well as extensively photographed in non-pornographic settings (Taylor and Quayle 2003, 6). These latter images should be treated as pornography, given that, as Taylor and Quayle note, they “complement and extend” the explicit pornographic material and provide contextual material about the child “making them more ‘real’ to the offender and fuelling sexual fantasies” (Taylor and Quayle 2003, 6).

²⁸ It is interesting to note that sometimes removing content to make it appear less pornographic can convey the idea that the editors believe it will be treated as pornography, leading it to in fact seem *more* pornographic. This can be seen in the case of *The Wall Street Journal* when they published Sally Mann's photograph of her four-year-old daughter Virginia in the nude but put black bars over her eyes, nipples, and vulva. In response, Mann claimed that “the censorship, not the picture itself, gave the image a tinge of pornography” (<https://www.artsy.net/article/artsy-editorial-sally-mann-s-photographs-children-viewers-uncomfortable>)

5 Conclusion

This paper began with the following question: what, if anything, is wrong with deepfake pornography such that the depicted person has a specific grievance against its production? The answer, I have argued, lies in the fact that there are two ways in which an image can be ‘of us’: an image can look like us and we can have a material connection to the image. Both connections can exist even if the image in question is a deepfake. And whenever an image is of us and is pornography, our consent is required. Pornography can be created by the *addition* of content or context, including the modification of an image such that it is now reasonable to believe it will be treated as pornography: deepfake pornography does just this by compositing someone’s image with underlying sexually explicit material.

Thus, the creation of sexual images that involve the use of someone’s likeness or photographs of them gives rise to a special demand of consent. Children cannot consent to be used—or to have their likeness or their photographs used—to create pornography. Thus, the production of deepfake child pornography—given that it starts with a photograph of a child—*necessarily* morally wrongs the children depicted. Deepfake adult pornography is objectionable whenever those whose image or likeness is used do not give consent.

My account does vital explanatory work. It explains what is wrong with *creating* deepfake pornography, independent of the wrongs of using or consuming pornography. It does not rely on any empirical claim about the causal connection between pornography and harm. And most importantly, unlike other arguments against pornography, it explains why the person whose image is used has reason to object.

5.1 Two Objections

I turn now to two objections that might be raised. One is that we *do* often consent to giving rights over—and therefore control of—the images we upload to various websites. It is true: few if any of us refuse to click ‘accept’ on the terms and conditions when trying to access a free service or website. Doesn’t my argument thereby imply that those to whom we have consented are perfectly within their rights to use our images, including to produce pornography? There are two responses to be made to this.

The first is that the standard of ‘consent’ when it comes to End User Agreements is notoriously poor (Barocas and Nissenbaum 2014). It often fails to meet almost any robust theory of consent (for example, being informed enough about the consequences and significance of agreeing, the practical ability to refuse when these services are vital infrastructure, and so on) (McDonald and FaithCranor 2008; Bakos, Marotta-Wurgler, and Trossen 2009). It can, therefore, only be called consent in inverted commas.

Secondly, consent in the case of sex and sexual images doesn’t necessarily function as other instances of consent do. Consider, for example, that in certain circumstances consent can be given by a third-party: if you are unconscious (or otherwise unable to make decisions), it is permissible for me as your nominated proxy to consent on your behalf about your medical care. However, it is impermissible for me to consent on your behalf to a sexual encounter or to the creation of pornography of you, even if you are similarly incapacitated. It is beyond the scope of this paper to explore further the special features of consent when it comes to sex and sexual images. However, the inapplicability of third-party consent in this area explains why a parent cannot consent to someone creating pornography of their child even though parents can consent to other things on behalf of their child, such as invasive medical procedures. It also gives some strength to the claim that, while we can consent online to many things,

we arguably cannot consent in the form of a blank cheque when it comes to sex or sexual images.

Another objection might arise from the seeming equivalency my argument makes between the wrongs of deepfake child pornography and deepfake adult pornography. My argument seems unable to explain how deepfake child pornography wrongs depicted children *qua* children. Here are two responses. Firstly, adults can consent and therefore there are (at least theoretically) some cases of creating deepfake adult pornography that do not wrong the person depicted, whereas the creation of deepfake child pornography *necessarily* wrongs those depicted. Secondly, there are perhaps special obligations that we have towards children who are in a position of trust and dependence that makes violations due to a lack of consent worse in the case of children than of adults. Note, for example, that both the rape of a child and the rape of an adult are wrong and wrong for the same reason—the lack of consent—and yet the rape of a child seems *worse* than the rape of an adult, if we were inclined to compare such things. Finally, the idea that child pornography is worse than adult pornography perhaps draws on arguments beyond concerns with the wrong done to those depicted: that the consumption and enjoyment of child pornography is wrong; that the causal arguments of harm to non-depicted persons strike many as more plausible in the case of child pornography; that the problems of inequality are more common or plausible in the case of child pornography. I do not disagree. However, this paper had a narrower goal: to establish how using someone's image to create deepfakes wrongs the person whose image was used. Thus, without undermining that what is wrong in both cases is a lack of consent, it is possible to accommodate the intuition that the creation of deepfake child pornography seems morally worse than the creation of deepfake adult pornography, even in cases where the adult depicted does not consent.

5.2 Beyond Deepfake Pornography

My argument has primarily focused on the issue of deepfake pornography, a phenomenon that is increasingly commonplace and therefore increasingly in need to philosophical attention. However, my argument goes beyond deepfake pornography, grounding out concerns about other kinds of activities and other kinds of images.

My account can be expanded to explain, for example, the wrongs of revenge porn, once we understand this phenomenon as the *distributing* of a sexual image of someone without their consent, noting that the distribution of pornography is another point at which consent is required, beyond the act depicted and the taking of the photograph. My argument can also explain how the creation of purely synthetic pornography (i.e. that does not start with a *photograph* of an individual) that intends to render an individual identifiable can be objectionable when that individual has not or cannot consent to the production of that image. Finally, my argument also has application to cases where the content, context or use of an image demands consent for their production. Pornography is one such case. However, there are likely to be others, for example using someone's image to advertise a product (Prosser 1960, 385) or to illustrate a story on, say, obesity ("Discussion Paper: Unauthorised Photographs on the Internet and Ancillary Privacy Issues" 2005, 11). This enables deeper reflection on another core case of deepfakes: the use of images of political figures to make fake political speeches or statements. Of course, there are arguments analogous to those concerning pornography that the creation of political deepfakes is wrong because the consumption and use of these images harms non-depicted persons. However, my argument can explain how *those depicted* have reason to object to their image being

used to make political speeches and statements beyond the political and epistemic impact on non-depicted persons.

My discussion of doctored pornography has implications for ethics and technology more generally: it establishes that there is a deep ethical significance of using images of real people for certain purposes that is not reliant on the truthfulness of the depiction. Fictions can wrong real people. The link between the depictions and the depicted—between a representation and the represented—in a time increasingly dominated by representations, can ground the interests of specific individuals. This paper lays the foundations to explicate how this interest is threatened, and those depicted wronged, even when they are not necessarily harmed.

References

- Adelman, Ronald W. 1996. "The Constitutionality of Congressional Efforts to Ban Computer-Generated Child Pornography: A First Amendment Assessment of S. 1237, 14 J." *The John Marshall Journal of Information Technology & Privacy Law* 14 (3 *Journal of Computer & Information Law*): 483–92.
- Ajder, Henry, Giorgio Patrini, Francesco Cavalli, and Laurence Cullen. 2019. "The State of Deepfakes: Landscape, Threats, and Impact."
- Andrea Dworkin. 1981. *Pornography: Men Possessing Women*. London: The Women's Press.
- Archard, David. 1998. *Sexual Consent*. Westview Press.
- Armagh, Daniel S. 2002. "Virtual Child Pornography: Criminal Conduct or Protected Speech?" *Cardozo Law Review* 23 (6): 1993–2010.
- Ashcroft vs. The Free Speech Coalition. 2002.
- Bakos, Yannis, Hlorenca Marotta-Wurgler, and David R. Trossen. 2009. "Does Anyone Read the Fine Print? Testing a Law and Economics Approach to Standard Form Contracts." *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.1443256>.
- Barocas, Solon, and Helen Nissenbaum. 2014. "Big Data's End Run around Anonymity and Consent." In *Privacy, Big Data, and the Public Good: Frameworks for Engagement*, edited by Julia Lane, Victoria Stodden, Stefan Bender, and Helen Nissenbaum, 44–75. New York: Cambridge University Press. <https://doi.org/10.1017/CBO9781107590205.004>.
- Benn, Claire. 2019. "Child Pornography in the Digital Age: A Conceptual Muddle." In *Pornography: Interdisciplinary Perspectives*, edited by Frank Jacob. Peter Lang Publishing, Inc.
- Benovsky, J. 2016. "Depiction and Imagination." *SATS* 17 (1): 61–80.

- Bergelt, Kelley. 2003. "Simulation by Stimulation: Is There Really Any Difference Between Actual and Virtual Child Pornography? The Supreme Court Gives Child Pornographers a New Vehicle for Satisfaction." *Capital University Law Review* 31:565–95.
- Brey, Philip. 2008. "Virtual Reality and Computer Simulation." In *Handbook of Information and Computer Ethics*, edited by K. Himma and H. Tavani. John Wiley & Sons.
- Burke, Debra D. 1997. "The Criminalization of Virtual Child Pornography: A Constitutional Question." *Harvard Journal on Legislation* 34:439–73.
- Carlson, Matthew. 2021. "Skepticism and the Digital Information Environment." *SATS* 22 (2): 149–67. <https://doi.org/10.1515/sats-2021-0008>.
- Chesney, R., and D. Citron. 2019. "Deepfakes and the New Disinformation War: The Coming Age of Post-Truth Geopolitics." *Foreign Affairs* 98:147–55.
- Child Pornography Prevention Act of 1996 (CPPA)*. 1996. U.S.A.
- Currie, Gregory. 1999. "Visible Traces: Documentary and the Contents of Photographs." *The Journal of Aesthetics and Art Criticism* 57 (3): 285–97.
- "Discussion Paper: Unauthorised Photographs on the Internet and Ancilliary Privacy Issues." 2005. Australia: Standing Committee of Attorneys-General.
- Doane, M. A. 2007. "Indexicality: Trace and Sign: Introduction." *Differences* 18 (1): 1–6. <https://doi.org/10.1215/10407391-2006-020>.
- Dworkin, Ronald. 1985. "Do We Have a Right to Pornography?" In *A Matter of Principle*. Harvard: Harvard University Press.
- Eaton, Ann W. 2007. "A Sensible Antiporn Feminism." *Ethics* 117 (4): 674–715.

- Eneman, Marie, Alisdair A. Gillespie, and Bernd Carsten Stahl. 2009. "Criminalising Fantasies: The Regulation of Virtual Child Pornography." *Proceedings of the 17th European Conference on Information Systems*.
- Fallis, Don. 2020. "The Epistemic Threat of Deepfakes." *Philosophy & Technology*.
<https://doi.org/https://doi.org/10.1007/s13347-020-00419-2>.
- Farid, Hany. 2004. "Creating and Detecting Doctored and Virtual Images: Implications to The Child Pornography Prevention Act." <<http://www.cs.dartmouth.edu/reports/TR2004-518.pdf>>.
- Floridi, Luciano. 2018. "Artificial Intelligence, Deepfakes and a Future of Ectypes." *Philosophy & Technology* 31:317–21.
- Gouge, Thomas A. 1965. "Peirce's Index." *Transactions of the Charles S. Peirce Society* 1 (2): 52–70. <http://www.jstor.org/stable/40319509>.
- Hansen, Jonas Kjeldbjerg, and Thomas Balslev. 2009. "Hand Activities in Infantile Masturbation: A Video Analysis of 13 Cases." *European Journal of Paediatric Neurology* 13 (6): 508–10. <https://doi.org/10.1016/j.ejpn.2008.10.007>.
- Healy, Margaret A. (ECPAT). 1997. "Child Pornography: An International Perspective." Sweden. <http://www.crime-research.org/articles/536/>.
- Henkin, Louis. 1963. "Morals and the Constitution: The Sin of Obscenity." *Columbia Law Review* 63 (3): 391–414.
- Hopkins, Robert. 2012. "Factive Pictorial Experience: What's Special about Photographs?" *Noûs* 46 (4): 709–31.
- Introna, Lucas D. 1997. "Privacy and the Computer: Why We Need Privacy in the Information Society." *Metaphilosophy* 28 (3): 259–75.
<https://doi.org/10.1111/1467-9973.00055>.

- Johnson, Robert. 1898. *The Complete Treatise on the Art of Retouching Photographic Negatives and Clear Directions How to Finish & Colour Photographs*. London: Marion & Co.
- Kauser, Kitty. 2007. *Shadow Sites: Photography, Archaeology and the British Landscape 1927-1955*. Oxford: Oxford University Press.
- Koppelman, Andrew. 2005. "Does Obscenity Cause Moral Harm?" *Columbia Law Review* 105 (5): 1635–79.
- Krone, Tony. 2004. "A Typology of Online Child Pornography Offending." <http://www.aic.gov.au>.
- Langton, Rae. 1993. "Speech Acts and Unspeakable Acts." *Philosophy and Public Affairs* 22 (4): 293–330.
- Levy, Neil. 2002. "Virtual Child Pornography: The Eroticization of Inequality." *Ethics and Information Technology* 4:319–23.
- Longino, Helen. 1980. "What Is Pornography?" In *Take Back the Night: Women on Pornography*, edited by Laura Lederer, 40–54. New York: William Morrow.
- Lovelace, Linda, and Mike McGrady. 1980. *Ordeal*. Secaucus, N.J.: Citadel Press.
- Mackinnon, Catharine. 1987. "Not a Moral Issue." In *Feminism Unmodified*, 146–62. Cambridge, Mass.: Harvard University Press.
- MacKinnon, Catharine A., and Andrea Dworkin. 1997. "In Harm's Way: The Pornography Civil Rights Hearings." <https://www.hup.harvard.edu/books/9780674445796>.
- May, Larry, and Marilyn Friedman. 1985. "Harming Women as a Group." *Social Theory and Practice* 11:207–34.

- McDonald, Aleecia M., and Lorrie FaithCranor. 2008. "The Cost of Reading Privacy Policies." *I/S: A Journal of Law and Policy for the Information Society* 4 (3): 540–565.
- Mcmullan, John. 2011. "The Digital Moving Image: Revising Indexicality and Transparency." *IM 7: Diegetic Life Forms II Conference Proceedings*, 1–16. <https://doi.org/10.1007/s10676-011-09522-1>. ISSN 1833-0533.
- NAPAC. 2016. "Media Guidelines for Reporting Child Abuse."
- New York v. Ferber. 1982, 81–55.
- Öhman, Carl. 2020. "Introducing the Pervert's Dilemma: A Contribution to the Critique of Deepfake Pornography." *Ethics and Information Technology* 22 (2): 133–40. <https://doi.org/10.1007/s10676-019-09522-1>.
- Patridge, Stephanie L. 2013. "Pornography, Ethics, and Video Games." *Ethics and Information Technology* 15 (1): 25–34. <https://doi.org/10.1007/s10676-012-9310-1>.
- Peirce, Charles Sanders. 1931. *The Collected Papers of Charles Sanders Peirce*. Edited by Charles Hartshorne and Paul Weiss. Cambridge M.A.: Harvard University Press.
- Pettersson, Mikael. 2011. "Depictive Traces: On the Phenomenology of Photography." *Journal of Aesthetics and Art Criticism* 69 (2): 185–96.
- Phillips, Dawn M. 2009. "Photography and Causation: Responding to Scruton's Scepticism." *British Journal of Aesthetics* 49 (4): 327–40. <https://doi.org/10.1093/aesthj/ayp036>.
- Poremba, Cynthia (Concordia University). 2011. "Real/Unreal: Crafting Actuality in the Documentary Videogame." PhD dissertation, Concordia University.

- Prosser, William L. 1960. "Privacy." *California Law Review* 48 (3): 383–423.
<https://doi.org/10.15779/Z383J3C>.
- Putnam, Hilary. 1979. *Reason, Truth and History*. Cambridge: Harvard University Press.
- Rea, Michael C. 2001. "What Is Pornography?" *Noûs* 35 (1): 118–45.
<https://doi.org/10.1111/0029-4624.00290>.
- Regina v. Hicklin, L. R. 3 Q. B. 1868, 360 371.
- Rini, Regina. 2020. "Deepfakes and the Epistemic Backstop." *Philosophers' Imprint* 20 (24): 1–16.
- Rini, Regina, and Leah Cohen. 2022. "Deepfakes, Deep Harms." *Journal of Ethics and Social Philosophy* 22 (2): 143–61. <https://doi.org/10.26556/jesp.v22i2.1628>.
- Ruiter, Adrienne de. 2021. "The Distinct Wrong of Deepfakes." *Philosophy and Technology* 34 (4): 1311–32. <https://doi.org/10.1007/s13347-021-00459-2>.
- Schriever, J. B. 1909. *Complete Self-Instructing Library of Practical Photography, Volume X: Negative Retouching, Etching and Modeling*. Vol. X: Negativ. Pennsylvania: American School of Art and Photography.
- Scruton, Roger. 1981. "Photography and Representation." *Critical Enquiry* 7 (3): 577–603.
- Shallit, Jeffrey. 1996. "Public Networks and Censorship." In *High Noon on the Electronic Frontier*, edited by Peter Ludlow, 275–89. Cambridge, MA.: MIT Press.
- Sontag, Susan. 2005. *On Photography*. Vasa. Rosetta Books.
- Strikwerda, Litska. 2011. "Virtual Child Pornography: Why Images Do Harm from a Moral Perspective." In *Trust and Virtual Worlds*, edited by Charles Ess and May Thoresth, 139–61. New York: Peter Lang Publishing, Inc.

Taylor, Max, Gemma Holland, and Ethel Quayle. 2001. "Typology of Paedophile Picture Collections." *The Police Journal* 74:97–107.

Taylor, Max, and Ethel Quayle. 2003. *Child Pornography: An Internet Crime*. Brunner-Routledge.

"The UN Convention of the Rights of the Child." 2010. The Children's Rights Alliance.

Walton, Kendall L. 1984. "Transparent Pictures: On the Nature of Photographic Realism." *Critical Inquiry* 11 (2): 246–77.

Williams, B. 1981. *Obscenity and Film Censorship: An Abridgement of the Williams Report*. Edited by B. Williams. New York: Cambridge University Press.