The Proper Likeness and the Models that Matter

Dr Bernie Hogan
Senior Research Fellow, Oxford Internet Institute
University of Oxford

For Channels of Digital Scholarship. Maison Française, Oxford. February 2, 2023

Release draft. February 2, 2023.

Abstract

The likeness can now be effectively encoded and used in order to render now synthetic versions. In light of this, we should revisit the notion of what constitutes a real likeness as opposed to a synthetic or fictional one. To do this, I draw upon the philosophical discussion about what is a real name, or more formally, what is a 'proper name'. I contend that we can elevate our discussion of deep fakes and synthetic imagery by considering the extent to which the imagery invokes a historically understood being. I use the notion of encoding and decoding from both information theory and semiotics to describe the encoding and decoding process. I conclude that a proper likeness is an analytically worthy concept for analysis apart from (or extended beyond) its specific instantiation in facial recognition. *This is a draft paper for a conference and likely to be integrated into a future monograph on identity*.

Introduction

The purpose of this talk is to provide a theoretical basis for the acknowledgement of the proper likeness as a meaningful and workable philosophical construct. The notion of a proper likeness is drawn in name and spirit from the notion of a proper name.

The reason for focusing on the proper likeness at this present moment is to find a more effective way to understand and respond to the challenges that AI systems provide to our notion of identity and our ability to wield identities authentically or otherwise. At present, I assert that discussions about computer-generated imagery and text are somewhat impoverished by the lack of meaningful acknowledgement of the proper likeness. This means that many issues are potentially wrongfully conflated or misunderstood. This misunderstanding partially stems from a naivete about the technology (either in terms of thinking it is either too stupid or too clever) and partially from a historical path-dependence which has blinded us to some emerging issues with social identity that technology of this form will inevitably confront.

Drawing upon the proper name

The notion of a proper likeness comes from the prior notion of a proper name. These names are understood by people to signify a particular person, object, or place. The study of proper names is onomastics, though the notion of proper names have been investigated in philosophy for centuries with significant contributions in thinking of them coming from Mill, Locke, Whitehead, Searle, and most importantly Kripke (1972).

Kripke challenges the notion of a proper name as being a marker of features of person. We might call that 'inverted duck theory'. Duck theory is loosely, if it looks like a duck and quacks like a duck, we call it a duck. Inverted duck theory (as I'm calling it) is more akin to Searle's notion of the proper name is the means by which we consolidate measurements about the self. We have things that waddle and quack - so we call them ducks. The problem with duck theory is its fragile dependence on counterfactuals. What if we have a duck that lost its voice or broke its poor flipper. Is it still a duck? Of course it is! Kripke here veers then towards the discussion of particulars. For example, if Obama was not elected president, he would still be Obama. Thus, inverted duck theory applied to people doesn't really work. We don't label him Obama because of what he does. He's labelled Obama because he needs a name in order for us to understand who we are talking about.

Kripke then asserts that the name really is a means by which we *fix a reference*. We use a name to fix in our communication (and in our reasoning) specific actors or objects which we believe have a particular significance. These objects are referents (or designatum in Kripke's vernacular). Although not herein, I assert the designatum bears striking resemblance to Benajmin's aura of artistic objects and that his classic Work of Art in the Age of its Mechanical Reproduction (1967). That essay foretold of necessary distinction between a likeness and a proper likeness by describing pre-technological art as having a historicity as an aura that was destroyed or obscured in the shaping of art as reproducible.

Considering the likeness as distinct from a name.

Drawing from Goffman (1959), social identity markers can be classed broadly into features that we give off and features that we give out. A proper name is a feature that is given out (or given to). It is the sense of give as a transitive verb. We tell people our name or we name others. A proper name is thus the means of fixing a reference that we intentionally communicate to others. A likeness is a means of recognition from that which we give off. We can shape our likeness in the backstage, selecting what clothes or makeup to wear, how to fix our hair or otherwise alter our appearance. But in the performance of the self, that appearance is not withheld from others intentionally. We inevitably give off an impression of the self (though our voice, our image to others, and our scent though culturally we do not usually attend so much to the nuances of scent). It is to some extent under our control (to remind, as is our name, too).

This meant that for the longest time, names were used to fix a reference to an individual because encoding the likeness directly was too difficult. Artists could paint a likeness and signal that it was who they say, but we could not institutionalise the likeness until we were able to record images, video, and audio. When this happened, however, our interest in the likeness as a signifier of the subject receded to our interest in the image itself as a signifier.

The likeness was thus heretofore only the concern of those who would profit from their likeness as a public matter. Royalty would worry about their image in a portrait, celebrities would eschew the paparazzi who would juxtapose candids with the curated likeness from the red carpet. However, in this shift we took the likeness for granted, partially obscured by the utility of the name. It is not the name that anchors the subject in tabloids but the recognition of the subject associated with that name. People do not pay for verbal gossip the way they do for images.

We may say that we have an impression of the likeness of other people that varies in its granularity but successfully allows us to recognise people directly. This likeness can allow us to identify both people in general and people in particular. While we may say it allows us to identify people in general, we really mean it allows us to identify first whether the object we are looking at is a person or not, but also whether it allows us to identify a specific person or not.

In the space of computer vision, we can identify thus two distinct tasks: likeness detection and likeness recognition. We normally associate this with facial detection and facial recognition. Facial detection is a means by which the likeness of a face is identified through coding a photo for certain distinct features and comparing those to an abstraction for 'is this a face'. The distinct features themselves can be arranged in such a way that we can generally detect a face, but also in such a way that the measurements for any given face are seen as distinct enough such that if one has an existing set of measurements for specific faces then a program can detect whether this particular set of measurements is close (or sufficiently close) to a set of measurements for a specific person. If the match is high enough, we do not simply have facial detection, but also facial recognition.

Articulating a provisional basis for the Proper Likeness

Facial recognition algorithms do not yet resemble the proper likeness but they are close. They are missing one key ingredient that has heretofore relegated the proper likeness to be an object of privacy rather than one of utility and exchange: a way to represent the proper likeness as a likeness. A facial detection algorithm on a camera might be able to detect my face a lot easier than distinguishing my face from another. Meanwhile, a street artist might be able to create an exaggerated caricature of someone that has proportions that would never been seen in real life and yet still manage to draw something that is identifiably comparable to the subject. So in the former case, we have an encoding of the identity that is then useable by systems which do encoding. In the latter case, we have a decoding of identity that is recognisable by the subject who is encoded.

But what about having a system that does both? We can envision a system that does this and as we have already seen, such a system is increasingly becoming a reality. This reality does not merely invoke the proper likeness. It helps us come to an understanding of what such a likeness is. That is, the proper likeness is premised on a means of both encoding and decoding the self that is given off in such a way that the decoded object could be recognised as the original subject. The proper likeness then is not a single object the way a proper name is (or is often considered), but it still has a single referent the way the proper name does.

So this means that some forms of recognition will be a likeness and a form of personal identification but would not feasibly be considered a proper likeness. For example, we may be able to uniquely detect someone from their typing speed. We may then be able to accurately track when people are typing at different computers (perhaps even anonymously given enough data). But then writing something with the accurate typing speed of someone else might not reliably allow a third party to state "yes, that's Bernie's typing alright" or do so confidently distinguishing between people.

One might be interested in whether they are identifiable from the information they give off. Given enough measurements and any two people (or objects) will be definitely distinguished from each other. This is a different albeit related matter related to the math of entropy. The matter at hand is whether the measurements can be used to reconstruct a novel representation of one's likeness (in a photo, video, or audio) which other people would remark bears as sufficient and unique resemblance to the original subject to indicate that the likeness has been involved. Can it be distinctly *recognised* as the subject. Can it be recognised by the subject themselves?

Returning to Dreambooth, AI Art, and operationalising the proper likeness at home

Perhaps the reason that this notion of a likeness has not been fully theorised is that we have not had the technology to encode the likeness in this manner for very long. We have had portraiture for a long time and we have had media technology for a long time. But in between these two is the realm of mediated portraiture, a field that only recently has come to be.

In art, Morris has recently introduced the notion of the real-likeness (Morris, 2020). This is akin to the proper likeness as stated here but from a different direction. In this book, in my interpretation, he acknowledges that a likeness is a form of a sign that links the signifier (which is the artwork) and the signified (the referent, meaning the person). While the book persuasively asserts the ontological status of the likeness, it does not do so practically, only by inference. The inference comes from the argument that we can make sense of a referent in a painting or a photograph and that this sensemaking only occurs because we are aware of and have been previously exposed to other likenesses of the referent. We need not even have seen the original referent (truly as no one today can claim to know the face of Christ but we can acknowledge a picture of Jesus as a picture of Jesus).

Part of the problem of treating this as a matter for the domain of art rather than science is that it creates a sort of black box where cognition resides from within which subjective experience of the likeness governs our acknowledgement of it. Yet, if we have enough photographs of someone we should certainly be able to say if a new photograph is of the same person. The new photo might have different levels of focus, lighting conditions, zoom, etc., and yet we can still recognise the figure. Further, computers can do this too. By training on and accommodating for a variety of lighting conditions and perspectives, we can algorithmically recognise faces. But more than that, once the algorithm to recognise the

faces is set, what's phenomenal is that the very same encoding can be used to generate new faces that resemble the original likeness.

They say impression is the highest form of flattery because it captures one's essence in the impression. In this sense, impressions are the means by which we distinguish between a likeness and a proper likeness. The proper likeness should be able to generate a representation (i.e. a photo, video, or audio) that sufficiently resembles the original subject to note that the representation would not have looked that way in the absence of training on that subject in particular. Further, the proper likeness should be invoked rather than circumstantial. With enough guesses, I could guess the names of most people. But that would not be invoking a proper name, would it Charlie Brown? It would just be coming up with names that look like proper names.

The website thesepeopledontexist is an example of a likeness generator. It generates photos of people that look like real people. Done enough times, there is a reasonable likelihood that it can be made to generate a figure that looks similar enough to anyone. However, the number of parameters required to render a face that uniquely signals a specific person suggests that we would need a very long time to come up with most people's faces by coincidence, much like guessing someone's name from scratch would be rather tricky.

In contrast to thesepeopledonotexist, the software Dreambooth is a proper likeness generator. Using Dreambooth, someone can use a small number of photos with a person's proper likeness to create a model that will allow one to render their likeness in a variety of settings. The proper likeness then exists as an abstraction in the model. At present, there is no single small file that includes a specific likeness that can be loaded into an existing model. Training models updates the weights of all the parameters in the models at present. Thus, the models (i.e. checkpoint files or large matrices of tags and their weights) are still several gigabytes large.

How to train a model for a proper likeness

The first thing to appreciate is that a face or a likeness needs to be encoded. Older models would measure the distance between key features like eyes inside a matrix of 'eigen faces' (Pentland, Moghaddam, and Starner, 1994). These measurements are generally pretty distinct and stable such that measuring these points on my face could create a wireframe that would allow one to distinguish my face from others. The problem is that they are very fragile to variations in perspective or lighting.

More recent models are far more abstract and inductive. There is no longer a specific wire-frame of 42 points or 42,000 points. Rather, there is only training and some forms of induction. Stable Diffusion is exemplary here. This is a generative model for images that is trained on billions of images. Each one of these images is tagged in some manner and the model reads the images and the tags, infers similarities (or complementarities), then creates a checkpoint file (or more recently a comparable tensorfile). This is a very large set of weights for everything from colour saturation to shape to subject. Running a generative

model one receives some output that is the statistical product of all the images that went into it. So you can ask for an apple on a chair underwater and it will produce an image inferring all of these things together. But it needs "apple", "chair", and "underwater". It will 'denoise' the model in such a way that the resulting product should have these things.

If you give it the term "Bernie Hogan" it will not know what to do as it has not been trained on my likeness specifically. However, if you give it the name "Barack Obama" then it will be able to render a figure that has features of Obama. However, this is the base stable diffusion model. This model is only as good as the data and the tags of that data. At present there are many discussions occurring online about what constitutes a legitimate source of training data for such models and whether the well tagged but copyrighted images from stock photos are legitimate sources of training data.

The important matter is that this base model can then be retrained. Dreambooth in the Stable Diffusion case is the software that can be used to retrain a model using photos of a subject. These then become a new tag in the model. So instead of 'model represented by checkpoint file' and 'tags' for navigating that model, we have $model_{prime}$ represented by a new checkpoint file and $tags_{+n}$. It is n and not just +1 since the extension can be run on multiple tags concurrently. Then by running the model with the tag signifying that subject, we then get that subject's likeness interpreted within the resulting image.

So, we can now render my face on that apple, or me holding that apple, etc...The key matter is that the likeness that we see can be identified as me rather than just a generic face. That checkpoint file (read: model) contains within it a proper likeness which is manifested by decoding the model.

The purpose of considering a proper likeness as distinguished from a likeness generally is that as an object, it can itself receive a name. The name doesn't have to be a proper name. It's just a label. So for example, if you take 50 photos of me and train on Dreambooth under the tag 'bh_test3', then when you run the model, you'll need to say "an apple underwater with the face of bh_test3, resting on a chair". If you do not invoke bh_test3, the data from those 50 photos will have a tiny but real influence on the look of the face, but it will not be my face. If you do not have a tag 'bh_test3' then you might be able to recover my face from the data, but in amongst the millions of other faces in the data it will be washed out. Thus, any attempt to recover it solely using key word descriptions like "white, bearded, oxford professor" etc...will be frustrated.

The consequences of the proper likeness

Named as an object in its own right, we can now think that the proper likeness ought to be considered as deriving from a specific identifiable (read:recognisable) person or object. It is not the object itself, but the encoding of that object that can be used to render a human-perceivable representation of that object. The proper likeness has heretofore not been given the philosophical status it is due because the encoding and decoding processes have heretofore been decoupled. We have previously encoded representations of people that bear their likeness but in doing so we focused on the media itself and not what the media signifies. We normally leave it to the mind to make the link between the image and the

caption. Similarly we have previously decoded the cognitively-encoded likeness but done this as a matter of artistic practice. The artist then becomes the means of encoding and decoding the subject. But again, the actual likeness itself resides in some shared cognition. We might say that it is 'intersubjective', insofar as it is subjectively established but done in a way that allows for a shared understanding.

Our brains are no longer the only devices that can encode and decode a likeness. (In fact, they never were the only devices; parrots can mimic voices with astounding accuracy, and octopi and chameleons will visually camouflage to look like specific environments; cf., Bender et al., 2021). The important matter is that the encoding process now can be connotative and exchangeable.

Denotative and connotative come from Barthes' Myth today as being first and second order signifiers (1973). To denote something is to represent it directly. I can point to a picture of a flag and say "That is a British Flag", thereby denoting the British Flag. I can point to a picture of eggs, bacon, black pudding, beans, etc...and similarly denote that it is a meal. But in the second case, I can also point to it and say that it is very "British". It doesn't denote Britain - one could have that meal anywhere the ingredients are available. But the fry-up occupies a special place in British culture (for better or worse).

The denotative likeness is the simple sign where we point to a picture and say "this is Barack Obama". This is because media are themselves first order. They map a representation of reality on to some encoding. This might be the light coming into the camera or the audio moving on the tape. But to paint a portrait of Obama that is not based on a specific scene, one would have to capture the essence of "Obamaness". That essence was the black box that connotes Obama. We have not cracked open the black box of subjective experience. But we have built another black box that accomplishes a similar thing.

These new black boxes are tools. Like any tools, we will have questions about their provenance, their second order effects (i.e. their unintended consequences), and their legitimate use. Fire has been great for cooking, not so great for home safety. Similarly in this case, we have significant concerns about who gets to invoke a proper likeness and under what circumstances. However, in the absence of referring to the proper likeness explicitly, we have hit the snooze button on this core philosophical question and stayed instead at the first order denotative level. We are still imagining that the technologies are doing a fancy job of cutting the heads out of magazines and pasting them on other photos. That was Photoshop twenty years ago. What the algorithms are doing is much more akin to what an artist does - it finds the general patterns and specific measurements that faithfully connote the subject and then renders a representation using these measurements and the expectations of the output.

Considered in this way, many questions now come into greater focus. The goal of this talk is not to answer these questions. It is to set up scaffolding for a research agenda that discovers as much as produces the answers to these questions. Below I take only two related examples: pornography and trade.

Pornography is a form of media representation that has a special status in human culture as it is meant to arouse one's sexual energies. Sex is generally seen as a private matter and the consumption of pornography, while popular, is not discussed fully in public discourse. The reason is likely that those who consume it are still shamed socially thereby making the conversations generally one-sided about the harms of pornography. This is not to reweight the discussion except to establish that there exists both consumers of pornography and producers of it in large quantities.

The harms of sexual imagery concern both the production and the consumption of pornography. On the side of production, we generally ask if the persons involved gave consent or if they even have the capacity to given consent (either cognitively or legally). Legally children and adolescents are not seen as having the capacity to consent to the production of pornography until their 18th birthday, a milestone that was not universal until recently, but has become a broadly understood international convention even if the age of consent for sex still varies between jurisdictions. Others that cannot give consent are those who have diminished mental capacities. Others still may be able to give consent but that opportunity was deprived of them as a form of assault (such as upskirt photos, etc...). To that end, there is legitimate issue about provenance of what can go into a model to begin with.

A likeness can be produced (read: encoded) from one set of source material and applied to another set of parameters. This has always been the case with art. Over two millennia there is probably not a single situation in which some artist has not inserted Jesus. And throughout history, individuals have been drawn into pornography generally without their consent. For example, in the early 20th Century, "Tiajuana bibles" existed which were small bawdy cartoons that invoked the likeness of celebrities often in sexually charged situations (Adelman, 1997).[1] These were poorly drawn cartoons, relying on the artist's encoding and decoding of the subject to faithfully represent it. It was generally of such poor quality that without the use of a name, many of these people would not be recognisable to the average reader, but nonetheless, some people would have still been iconic. More sophisticated artistic albeit sexualised renderings have existed throughout history. For example, in 2020, Cree artist Kent Monkman created a storm in Canada for a highly controversial painting that depicted his drag persona Chief Eagle Testickle sodomising Prime Minister Justin Trudeau publicly while being mocked by First Nations women. The artist is speaking in some senses to the plight of the First Nations and critiquing Trudeau's enthusiasm as misplaced, insincere, and performative. But the image itself exhibits significant sexual violence and created sufficient questions to provoke an apology from the artist (Grabish, 2020).

In more recent years, the use of Photoshop has similarly been applied to porn. It originally did not have the means to 'encode' a likeness. It encoded denotative media (such as a photo) and then manipulated the state of the photo. This process can be very good, especially among professional graphic designers, but it is not the same as invoking the proper likeness. They might take one photo from one place, warp, shift, and crop it, alter

[1] Before wincing too hard at the invocation of Tiajuana Bibles, one might read Kripke's essay in Adelman (1997) about their linguistic creativity via anonymity. That's not Saul Kripke's essay, but is nonetheless of esteemed late librarian Madeline Kripke, and indeed, Saul's younger sister.

the lighting and such and pasted it into the new photo. The artist will this work with the photo-as-data but not the likeness as data.

It is through recent advances in ML that we have been able to finally recreate this encoding process (though it is not for me to say how much this process mimics or resembles the subjective encoding done by people). Thus, instead of taking a specific face from one magazine, cropping it, adjusting its size and lighting and pasting it on to another photo. The image is now rendered from a statistical distribution of likenesses (of things like torsos, backgrounds, lighting conditions, etc..) and a particular proper likeness of the person on whom the data has been trained.

So it is from here that we get deepfakes. Unfortunately, the name 'deepfake' is somewhat overloaded. Deep comes from deep learning (or the learning of statistical averages rather than specific details) and fake from the fact that the image is rendered. The name predates the current flood of AI-art and unfortunately has a sort of path dependence to it, since it could refer to any synthetic AI image, as these are all made up from the model (thus implying they are fake) and employ deep learning.

So here we can arrive at a more nuanced distinction: a deepfake refers to the creation of an image that invokes the likeness of an individual. To that end, not all synthetically generated porn would qualify as a deepfake. Rather, a vast array of such porn signifies a generic person, or a person according to the prompts (male, tall, handsome, etc...). However, deepfaked porn in my frame necessarily includes the encoding and decoding of a specific proper likeness.

We may then ask about two specific processes which themselves could lead to harm, but are independent. One is the encoding of a proper likeness. Indeed, every photo that a program trains on will included likenesses, by nature of having literally any subject in a frame. Insofar as we can detect objects or even colours we can say an image has a likeness. But training on images does two things - helps us recover what a coherent image looks like, and helps us recover what a coherent image of this object looks like. The first can train on likenesses without encoding a proper likeness (except incidentally, and to the extent that a specific proper likeness emerges we might say that the model overfits on a specific person). So we can train on 1000 faces without labelling the specific names of the faces. Labelling the images tells us what things we can decode from a model. If we label some of the faces as linking to a specific person and make that label available in the model, we have now encoded a proper likeness even if we named it "12modelFace2".

Synthetic pornography is here and will remain here. It is as old as art. However, deepfake pornography is then synthetic porn + proper likeness. Synthetic porn is not deepfaked by necessity except where it invokes an encoded proper likeness.

We might counter by saying that the model which renders the pornography was itself trained on one's likeness therefore it is a likeness. Yes, it is a likeness, but it is not necessarily a *proper likeness*. This rests on the distinction between derivative and transformative work. If one cannot explicitly recover the likeness from the model directly, then in my opinion the work is necessarily transformative, although many legal minds are

themselves considering this at present (including the now somewhat unpredictable US Supreme Court).

What this work adds is a grammar for the key concern in deepfaked pornography. It is not the presence of pornography which is at question. We understand as a society that people will photograph or video themselves without clothing and will do this for the arousal of themselves and others. However, we also understand that self-image is an important part of one's sense of self and self-integrity. To see oneself in a situation with verisimilitude that does not reflect that person's experience is upsetting, particularly when that image is used for the satisfaction of others. This then becomes a form of exploitation of one's proper likeness. Insofar as the sexual satisfaction comes from the titillation of that particular person and not a person of broadly similar measurements, it is not about the mere geometries of physical attractiveness but about the power dynamics that occur when our proper likeness is invoked (and especially without our consent).

This further adds some weight to the potential legal ramifications of training a model on one's likeness non-consensually. There exists considerable difference across the world about one-party recording and similarly we might want to extend that discussion into one-party training. If I can record someone (thereby creating an object that denotes their likeness) can I therefore generate a likeness from that recording (thereby creating an object that connotes their likeness)? If we have special exceptions from one-party recording for sexual imagery should we have it similarly for one-party encoding or decoding? What we can say is that the rendering of the likeness of a human form is not the same thing as the rendering of a specific human and that this insight ought to be propagated.

The second issue concerns trade. The decoding of a proper likeness used to be a service. That is someone who is trained in the creation of likenesses could be employed to create a likeness. This would be a portrait painter or photographer traditionally. That was because the encoding and decoding processes took individual skill. However, as these two practices become increasingly automated, the act of encoding and decoding become more about the access to the tools than one's skill in this matter (although discussions online suggest that 'prompt engineering' is itself becoming a form of artistic practice as one learns how different prompts alter the model in a more broad or holistic way). One who would generate an image needs access to the generative technologies (the encoding and decoding engines) but also to the object itself that has been encoded and can be decoded.

We will want to consider how the proper likeness, as a literal property of the self, should be considered in terms of trade. A likeness could be licensed in perpetuity or locked up in a contract. One might be able in some legal regimes to either sublicense one's proper likeness or invoke it using some platform logic. For example, I might be able to make deepfakes on a platform that include me and my friends, but not people I am not friends with. One might be able to store a model of oneself as rendered that can be licensed to a third party. As these systems continue to evolve, we will wonder how they are using the data for our likeness. Most people will not be aware of the likeness rights they sign away on EULA's but the cultural diffusion of the proper likeness may allow for a more nuanced conversation about what can be used by whom. Training on my face because it is a face might be fine for

most people, but training on my face because it is me and you want to render me should place it in a special class of practice.

A research agenda on the Proper Likeness

As a connotative rather than denotative object, we might ask how good does a likeness have to be? The follow-up to this is, good enough for what? If we say to fix a reference, we might ask 'for whom'? To whom does the referent need to be fixed and for what purpose?

These questions then lead us to a discussion of algorithmic fairness, either concerning bias in the data or governance of the trained model. They also lead us to a gray area where we might need to disambiguate an impression from a deep fake and fair use of parody from exploitation of the self.

Some of these questions will not be solved analytically. Our own cognitive faculties permit and inhibit some basic expectation of face detection or voice detection. Recall that I stated that a proper likeness needed to be able to establish for the viewer the unique resemblance to the referent. How this changes for celebrities, historical figures, strangers, and people from different cultures are all meaningful questions. Will such tools strengthen the social bonds we use to establish shared referents? Will they be deployed as political means to re-envision referents (as visualising "American President" in DallE might include women even though none have yet been elected)?

Now that we are on the other side of the uncanny valley can we see what it looks like before declaring it someone's (or some corporation's) territory?

To conclude the talk in relation to digital scholarship, this is not the end of a research agenda on things which have already been encoded. It is the beginning of a new era of research on what can be encoded by whom and for what purpose. With newfound interest in the proper likeness, we can move beyond seeing AI art as a threat and digital research as something that takes place with existing media towards seeing AI as a means of inference for referents that already make sense to us.

Our notion of digital research then can steer towards exploring how useful this new means can be and how we can govern its application as we use it to envision (sometimes literally) a new world.

References

Kripke, S. A. (1972). *Naming and necessity*. Blackwell Publishing.

Adelman, S. (1997). *Tijuana Bibles: Art and Wit in America's Forbidden Funnies*. Simon & Schuster. New York.

Barthes, R. (1973). Myth Today. *Mythologies*. Hill & Wang.

Bender et al., (2021) ‘On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?’ FAccT '21. Virtual Event, Canada.

Benjamin, W. (1967). The work of art in the age of mechanical reproduction. In *Illuminations* (pp. 217–251). Schocken Books.

Goffman, E. (1959). *The Presentation of Self in Everyday Life*. Anchor Books.

Grabish, A. (2020). "Artist Kent Monkman's painting of partially nude Trudeau with laughing women creates uproar online". CBC News.
Available online: https://www.cbc.ca/news/canada/manitoba/kent-monkman-backlash-trudeau-painting-1.5577452

Moriss, M. (2020). Real likenesses: Representation in Paintings, Photographs, and Novels. Oxford University Press.

Pentland, A., Moghaddam, B., & Starner, T. (1994). *View-Based and Modular Eigenspaces for Face Recognition*. Available online: https://citeseerx.ist.psu.edu/doc/10.1.1.47.3791