

### **An over-optimistic assessment of what we could accomplish as a class, at least over the next couple years.**

One of the places that digital humanities thrives is in the intersection of new algorithms with new resources. I came into this field as a scholar of books and texts, and will likely remain so. But it's increasingly clear to me that someone—and hopefully someone in this class—has an enormous amount to gain from working through the ways that **digital images** will be a useful source for digital history.

So as we look at the variety of ways that scholars have used texts, maps, and networks in the past, we'll be making some touchstones into thinking about we *could*, *can*, and *will* do with historical images and new image processing technologies.

I don't know as much about this as I do some other things I teach. But I think finding the holes in my knowledge—and learning on your end how to explore a rich and growing literature and think about applications for digital history—can be extremely useful.

### **Why digital images have a higher research upside than digital texts:**

There are several reasons I think that the research upside for digital images is, at the moment, higher than for digital texts.

1. **The algorithms have improved much more rapidly than for text, so fewer low-hanging fruit have been picked.** Five years ago, the landscape of text analysis didn't look especially different than it does now; there's been just one major breakthrough from text processing to come out of the neural network world (word2vec, which itself simply makes easier and better things you could already do). A few things historians don't use, like sentence parsing, seem to have gotten better; there clearly *will* be breakthroughs in OCR and named entity recognition, but they haven't filtered through in a way where we can easily use them right now. Images, on the other hand, have gone nuts; facial recognition is much better, the ability to identify objects from pictures suddenly exists, captioning works fairly well, and the new embeddings of images in high-dim space now seem to encode high-enough-level content that they can do completely different things than were possible before.
2. **The algorithms are better than for text analysis, full stop.** This is kind of contentious. You can't really know the answer to something like this! But basically, the things that researchers can do with images, and have been doing since 2012/2013 or so, have a much higher "wow" factor than what they can do with text. The reasons for this are both complicated and maybe opaque; we'll explore some of these things in the class, and at the end of this post. It's certainly the case that some vaguely analagous tasks, like object recognition for images vs subject classification for documents, and more successfully solved for images (especially since images aren't in multiple languages.)
3. **The pre-existing state of the art for algorithmic discoverability is worse.** Ted Underwood once wrote that "[Search is the killer app for text mining](#)", and that it infiltrated scholarly practices in the 1980s and 1990s without us quite noticing. Image search, on the other hand, has been basically non-existent. So we're just starting to be able to look at this vast archive. But rich contextual search on large image collections *is* possible. What would it look like? What would we want it to do?
4. **The pre-existing state of the art for library discoverability is worse.** Books are well-indexed by the Library of Congress and you can browse through stacks to find the important ones. Photographs

do often have metadata, but the form varies dramatically from institution to institution. Library of Congress Subject headings are often used with institution-specific themes; from the outside, it looks to me like the standards for what you label in an image are even less clear. One of these pictures is of the subject heading “trash,” one is “refuse;” poke around pictures of dumps, and you’ll also



see “refuse disposal”, “refuse and refuse disposal”, and a whole bunch of others.



There are often really good reasons that library catalogs make the precise distinctions that they do, and you should never be too quick to criticize them; but there are often local, specific, and other reasons that integrated search for images across large repositories is hard.

5. **The sources are fresher to the field in accessibility.** This is kind of corrolary to the last. The books in Google Books have been available to every scholar with a library card for decades; the images, on the other hand, have mostly been locked down in a single institution’s archives and so seen by, likely, many fewer people.
6. **The sources are fresher to the field in publishability.** One of the reasons that historians use visual sources less often is because they’re expensive to print. People get MacArthur grants and are like, “oh

great! Now I can put color photos in my books!” Web-oriented publishing is image-friendly; so on the output side of digital humanities, there are some great possibilities.

7. **They make for better visualizations.** This is a bonus, nowhere near as important as the others. But still; it’s one of the interesting differences between books and pictures that you can learn much more about a picture in half a second than you can about a book, just by looking at it. The web lives and dies by images.

## What can we do?

This is one of the big questions for the class. I hope we’ll be able to both generate some new ideas and start to test some of these out for plausibility. But in the month or two I’ve had this in the back of my mind, I can see several interesting possibilities.

The basic rule here is *feasibility*. I’m not trying to cook up science fiction scenarios below: I’m trying to cook up things that I could write in a week or less. That means, in practice;

1. There’s already a paper on ArXiv outlining similar tasks;
2. There’s code (preferably tensorflow) for doing it;
3. There’s an existing model to work off of. (For faces and image recognition, we can work from what’s known as the “bottleneck” layer of a prebuilt model, which means much less training is required.)
4. The math isn’t crazy hard. A lot of this relies on embeddings of images, which in terms of navigation is pretty much the same as embeddings of words, which some of us at Northeastern know a fair amount about and can help work the problems.
5. Computational feasibility. This is where I’m most worried at the moment; I find this stuff to be kind of slow compared to the text work I’m used to, and Research Computing hasn’t successfully gotten TensorFlow running on their monster GPU machines. If we come up with a cool idea, though, there a number of places we can go to, hat in hand.

So: what are some possibilities? This list may be expanded.

### The central face repository database.

Someone will build this, eventually. It will be really cool; you’ll read about in wired; it will be acquired by Facebook or [ancestry.com](http://ancestry.com) and I’ll grumble that I had the idea first and if only I had worked on it in time *I’d* be rich. We build a database of every face in the DPLA, vectorized the way that security services do it. So we’re starting with a lot of faces; gotta be at least 10 million, I think, which is possible.

Now: you upload any photo query and say: find my photographee in it! This is basically like the famous photograph of Adolf Hitler in the 1914 crowd celebrating the declaration of war. How did they find that photograph? I don’t know; either it took the full efforts of lots of Nazi party combing through pictures to spot Waldo in there, **or** they just airbrushed Hitler in later. Let’s assume the former; suddenly, neural networks are like a little Nazi party of your own you can use to comb through photos for your own genealogy. (We’ll tune up that sentence before the VC pitch).

It would work something like this: You upload your old family photos with labels on the people you know. You then can search the full database for all the times your grandma photobombed a wedding; and you can get fruitful leads on the unknown faces in your photos to discover who else has connections to your great-granddad’s union bowling league.



Figure 1: Hitler

Is this feasible? I don't know. A friend who's a computer vision researcher at Google told me at a party this weekend that it was a solved problem. On the other hand, OCR is also a "solved problem," and OCR is generally regarded by humanists (not necessarily me) as unacceptably bad. And clarifai, which is not just some niche player, [can't even find Waldo a lot of the time](#). So we'd have to see.

For historical search, it could also be quite interesting. Some subaltern people may not exist in written records; but they may drift, Zeliglike, through the historical record in the background.

And don't you want to know what previously undiscovered photos of the young Harry Truman lurk in local archives in Kansas City?

There's a cool social network analysis lab to be done here too; see who gets photographed with whom, and find the Kevin Bacon of the social scene. Just what photos are appropriate here is a good question.

### **The search for object X.**

This one's straightforward, but requires a little more handholding. Take, for instance, photos of Model T and Model A Fords; get enough that you can train a classifier that finds them relatively well. I think this might be possible, though I can't say for sure. (Based as much on my ignorance about Model T's as image analysis.) So you could look at nationwide pictures to see how commercial products spread; how late were model T's still showing up on the streets in large numbers? Did German cars disappear in 1917, or suddenly start showing up with their badges missing? What streets in the United States displayed the most NRA member signs in 1934? What were the most widespread representations of Jesus in photographed American homes in the 1920s? Etc.

None of these are amazing instances. But there might be a great one out there. (Advertising history seems especially fruitful).

Also, there's the possibility of interesting sources for combining object X with field Y. We know a lot about the demographics of computer science in the 80s and its relationship to hobbyist culture in the 70s can we look to see men and women in the pages of Byte Magazine, and see both whether they exist, and what kind of thing they're doing? Can we do this across all sorts of advertisements?

### **The critical digital approach: what *don't* computers see?**

Reading through Imagenet labeling results on DPLA lately, I've been freaked out by the way that it refuses to see human beings. There are no humans in the training set, so it sees them simply as some sort of strange haze on which the real things of the universe are displayed.

For instance, here's a DPLA photograph: 

And here are the most likely objects that the Inception network sees in it:

```
cowboy boot (score = 0.18367)
miniskirt, mini (score = 0.07054)
sock (score = 0.06981)
fur coat (score = 0.04957)
park bench (score = 0.03653)
crutch (score = 0.02502)
curly-coated retriever (score = 0.02416)
```

That's understandable, because those are imagenet labels, and this image is outside the training set. But it's also deeply spooky. It can't see people. Anyhow, it does see a lot of *clothing* which isn't there, but which is—for me—kind of summoned up by the rest of the image. That dude on the left *might* be wearing cowboy boots we can't see; that woman in the middle might have worn a miniskirt in the last week.

Running the straightforward Imagenet categories on 19th century objects is unlikely to be super-fruitful in terms of bulk-describing what's in them. (It tends to see a lot of webpages and magazines.) But it could be very useful in terms of understanding how and what computer vision is learning to see, and how that can be better understood.

In terms of *talking about* the digital today, the ability to think the way that computers is incredibly fruitful and still only vaguely understood.

### **A long Westworld Digression (sorry, class.)**

Long digression: I find the missing people in the computer vision labels spooky for reasons I can't fully articulate. Part of it has to do with the way that artificial "intelligence" is far more controllable than human intelligence in what it can be taught to ignore. There's an essay about Westworld that I'm never going to write, or which I might write as the introduction to the version of this approach, or that I might beg Aaron Bady over Twitter to write. It's about how the oft-repeated "that doesn't look like anything to me" line, in which the robots can't see photographs of the present day because they've been programmed not to, also appears to apply to the way they see black-white racial issues. That is wildly anachronistic for a show about the West, but entirely on point for what a 2060 western themed park would be, because obviously you can't have your black customers getting harassed by androids all the time. So it's probably true for legal reasons that those androids are like Stephen Colbert's old character and can't see race at all. This is not dystopian; but neither is it precisely utopian. It's totally technologically feasible, and in fact seems to represent how Silicon Valley would like to think about race; let's eliminate all prejudice, and ignore inequality that arises. I think that there's a major lesson to be learned here for the humanists and computer scientists who happily assume that datasets will just replicate existing inequalities. There's a science fiction world where government or market forces instead make computers spookily, weirdly color blind in ways that don't replicate existing inequalities but end up doing something much stranger.

### **End Westworld digression and back to critical data studies.**

There are a lot of extremely pressing questions about the ways machine learning replicates biases in the world; these become even more interesting when we throw them into the past, because we can potentially see in new and interesting ways the ways that they fail on older sources. And this could tell us two things:

1. It could help reinforce a critical data worldview that has pushed back at the objectivity of machine learning by showing how it instantiates biases. There may be all sorts of class, race, and gender ways that these fail on the past that are suggestive of ways they fail in the *present* that we might not see so critically.
2. It could help highlight the difference between the present and the past. Are there things in old images that present-day networks find wildly improbable? Can we find them? What are they? Do they look strange to us? Can we iterate to find really odd things in the past?

### **The neural-network-cartography nexus**

I've seen a little work here, but there could be more. Just in terms of presentation, I'd love to see a pix2pix network that takes maps from one atlas style and converts them to another. A good georectified database (which I think

David Rumsey, a DPLA contributor, has) could allow you to find multiple different representations of the same place; you could use these to build up a complicated underlying representation of each region with different visual manifestations, so you could even make maps that combine house cartographical styles in weird ways: that look like 2/3 Stieler and 1/3 Johnston, with a tiny dash of Cram. Maybe this wouldn't work.

Or maybe we could throw massive datasets of railroads shapefiles and maps at pix2pix, get real good at extracting railroads out of some style of map, and move on.

One obvious and important project is: use the extensive NYPL fire insurance maps and building footprint data as a pix2pix trainer, and then run off some transfer learning to start assembling raster data about buildings for the rest of the country that could then be vectorized.

More important still would be fully unsupervised neural network georectification. I wouldn't know where to start on this, but it seems at least possible that some Inception-based architecture could learn to match various distinctive stretches of global coastline and properly orient them. This is, I think, the closest to science fiction I can think of.

### **What *should* we do?**

This is a question with two parts. One is professional; is it a waste of time to count certain things because no historians, or members of the public, care about them? Then let's step away from it; one of the reasons we need historians is to temper the urge that lots of scientists, and plenty of humanists (myself too much included) sometimes feel that just because something *can* be counted, we *should* count it.

### **Ethics**

But the other side is far and away the most interesting, and in some ways presents the most important reflections for the humanities writ large, not just history. Mulling over what the data makes possible, I quickly encounter all sorts of ethical problems that everyone in the course will have equal standing to address. I'm going to be a little glib here because if ethics wasn't supposed to be funny, it wouldn't be a field predominantly concerned drawing cartoons of babies on train tracks about to get run over. If I cross any lines, please let me know.

A lot have to do with faces. Physiognomy has made a comeback in the last few years; it's probably as specious as ever. But one could easily build a classifier that would try to identify the race of a person based on their facial structure. You could use that to look for previously ignored African-Americans in settings where they were quite improbable; to track the slow integration of schools or social clubs after 1965 based on photographs of meetings; or any number of other things. But it's also, maybe, a deeply offensive idea; it relies for its efficacy on a biological notion of race that has been discredited in scholarly work,

Or consider efforts to search for individuals. Suppose that you build a classifier that goes through old lynching photographs, pulls every face it sees, and then tries to find near analogues of those faces elsewhere in the public record and builds a database that adds to the web records of all these 1920s figures that they were involved in one of the great forms of racial atrocity the country has ever perpetrated. Maybe that's OK; what sort of protection do (almost certainly dead) lynch observers deserve? But there's certainly a cost; we know, for instance, that Charley Guthrie, father of Woody Guthrie and grandfather of Arlo Guthrie, probably participated in the lynching of [Laura and L.D. Nelson](#) in 1911. That was a challenging thing for the Guthrie family descendants to deal with, and they at least were anti-lynching. What would it do to dig up a previously undiscovered photo of some Arkansas grandmom's beloved father participating a lynching and stick it on Google? I mean, I guess we can hope it would be a learning experience. But I know this one guy who found out last year that his father was

arrested at a Klan rally in 1927 in Queens, and he actually seems to have gone off the deep end in the *other direction* lately; he started reading some of those white supremacist sites, retweeting them. He even put one or two in his cabinet!

Another group claims to have developed a classifier that can tell gay from straight men by looking at their faces 81% of the time. This has got to be charlatan hogwash, right? I don't believe it's true. (It reminds me of the first line of the great Will Ferrel *Pearl the Landlord* sketch: "Well, we got the test results back—my dad's gay.") But suppose it isn't; should we start throwing it at Brahms and Lincoln and everyone else who anyone's had questions about, just to see? Or looking for group photographs of monasteries or print shops or labor union locals where the classifier thinks *everybody's* gay, and trying to figure out the potential implications for the process of social formation?

And what are the ethics of persons photographed against their will? The history of photography is rife with stunning examples of photographs being taken, [even of relatively high-status people](#), that should probably just be destroyed. (Not something I'd say about books.) There's a lot of fascinating ethnographic photography, for instance; in your typical app store land, a large batch of paired images like this image I took from the American



Yawp textbook

could make for an extremely inappropriate iPhone app, but in better hands it might produce a fascinating and well-attuned account of something else about what changes. But how much freedom did that 18-year-old have when he showed up at the school?

One interesting meta-question is: **can we rely on archivists' judgements about what's appropriate to scan and publicize, and what must remain private; or do the new uses just coming into view potentially demand a rethinking of appropriateness based on uses that could not have been imagined even ten years ago (accurate identification of unlabelled pictures, say) require a new rethinking of what's ethical.**

**Are ethics less important for dead people?**

We historians don't have to deal with institutional review boards because our subjects are dead.

Is it OK for us to do some really spooky stuff on dead people? Is it even useful, if it helps raise awareness of just how much the Chinese government probably knows about you, gentle reader, somewhere in its own databases?

The world of data processing often leads to some truly evil ideas. It's probably not the case that demonstrating their feasibility in the past is an unadulteratedly *good* thing; but it may have some benefits, especially when the dead subjects are not especially vulnerable themselves.