



BETHANY NOWVISKIE

≡ Menu

reconstitute the world

Posted on 12 JUNE 2018



*[What follows is the text of a talk I gave in two different contexts last week, as
"Reconstitute the World: Machine-Reading Archives of Mass Extinction." First, I*

opened the summer lecture series at the University of Virginia's Rare Book School, where I'm privileged to be a faculty member and supporter. Next, I closed the first week of the 2018 Digital Humanities Summer Institute (DHSI) at the University of Victoria and opened a Digital Library Federation (DLF) unconference on social justice and digital libraries, DLFxDHSI. I started my UVic talk by noting that we met on the unceded, traditional territory of the Lkwungen-speaking peoples of that part of the Pacific Northwest, and I therefore acknowledged the Songhees and Esquimalt, and also the WSÁNE? peoples who are among the First Nations with historical and enduring relationships to that land. I note this here, because the talk I gave is relevant, I think, to the need for humility, respect, and reparation and to the celebration of endurance and renewal (or, better, reclamation) that such statements, still uncommon in the United States, suggest.]



This is a talk on digital stewardship and *heritage futures* at a strange confluence. I'm more used to saying "cultural heritage"—*cultural* heritage futures—and I will certainly be addressing those today: possibilities for the strongly future-oriented digital stewardship of human expression as we encounter it in transitory, embodied performances, as intangible culture, and of course in ways that leave more lasting, material traces. But I use the broader phrase "heritage futures" deliberately, because this is also a talk that moves me beyond my training and my various cultural comfort zones in two big ways.

First, I'll step out of the humanities to gesture at projects in preservation, access, and scientific analysis that address our broader, global heritage of biodiversity. That's a heritage we share with all living things. And where we've failed in *stewarding* living environments, I think it's fair to say that we've only moderately well succeeded in *documenting* them—which in this case are two radically different things. Our success is particularly mixed—though improving—in documenting them with an eye toward the

activist, artistic, or reflective work we may soon wish to do in radically changed ecosystems.

Next, I'm here to speak, frankly, far beyond my own expertise, but I hope with some imagination, about how we might connect these concerns to our present revolution in machine learning and artificial intelligence. I particularly want to think about how to do so in a way that leverages the skills and deep-seated understandings that a background in the humanities, librarianship, or in post-custodial and community archives almost uniquely provides. It's important for me to say, though, that that there are some lenses or comfort zones that it is difficult for person coming from a settler background to drop and exit, particularly when talking about library and museum collections "acquired" and maintained in colonialist contexts. I'm trying.

I draw the title for this talk from [Adrienne Rich](#)—from part of a 1977 poem she called "Natural Resources."

*My heart is moved by all I cannot save:
so much has been destroyed*

*I have to cast my lot with those
who age after age, perversely,*

*with no extraordinary power,
reconstitute the world.*

Sounding through today's talk—alongside a bit of ecological despair that is still echoing, for me, from [my last attempt](#) to address these issues in front of a DH audience—you may hear the undertone of a [feminist ethic of care](#), and also of that utterly commonplace and yet counter-acting power of reconstitution or [repair](#) that Rich evokes in these lines on the screen.

But my basic argument is simple, and it has to do with *stuff*. What I want you to take away from this talk is an understanding that the *constitution*—the very make-up and organization—of our natural history and cultural heritage collections becomes vastly more important when we accept two truths. The first is that we assemble them at the end of things. All “[archives](#)” of the Holocene (and therefore not just of print and manuscript culture and their digital sequelae, but indeed our archaeological and more recent paleontological records, and the stories we read in landscapes and ice cores)—all are archives of diminishment: of a shift to plant, animal, and human monocultures. They are archives, in fact, of the 6th great mass extinction of life on our planet. And accompanying that sobering thought is a second necessary understanding. The very make-up (again, the contents, the structure) of our heritage collections likewise becomes a matter of critical concern, when we realize that we no longer steward them for human readers alone. *This* is the strange confluence of our present moment.

In the same way that human beings are shaped by what we read, hear, and see, the [machine readers](#) that follow us into—and perhaps beyond—the Anthropocene have begun to be molded by increasingly “unsupervised” encounters with our digital libraries. I’ll describe these encounters in more detail in a moment, as I ask you to consider what we offer up for study in our heritage collections now, and how we might better conceptualize archives of mass extinction for the living generations and artificial intelligences to come.

As I talk, I’ll invite you to dwell on some further questions. Some of them will be posed in a more implicit way, rather than explicitly in the talk, so I’m going to lay them out now. What kinds of indigenous knowledge do we neglect to represent—or fail to understand—in our digital libraries? What tacit and embodied understandings? What *animal* perspectives? What do we in fact choose, through those failures, to extinguish from history—and what does that mean at *this precise* cultural and technological moment? On the other hand, what sorts of records and recordable things *should* we let go—should we be working as hard as possible to *protect* from machine learning for the good of vulnerable communities and creatures—knowing, as we do, that

technologies of collection and analysis are by nature [tools of surveillance](#) and structures of extractive power? And, finally—from an elegiac archive, a library of endings, can we foster new kinds of human—or at the very least, humane—agency? This is a concept I’ve formerly called “[speculative collections](#),” which I’ll ask you to think of in the current context as digital archives from which to deep-dream new futures.

As Adrienne Rich suggests, the most ordinary (and still extraordinary) power we mortal beings possess is the power to make poetry from fragments of the past. We’ve begun to extend it, through machine learning, in uncharted ways. Might it be called on, one day, to reconstitute the world?

Cabinets of Curiosity

That was heavy. I’m going to start with a lighter overview of some relevant projects.

The [Biodiversity Heritage Library](#) is a Smithsonian-based international consortium and digitization collective of botanical and natural history libraries. It was the winner of DLF’s inaugural [Community/Capacity Award](#) in 2016 and leads the development of a lovely and effective “biodiversity commons,” complete with APIs for computational access of various kinds, harvesting and delivery through partnering nodes worldwide, full text search and faceted browsing, and important contributions both to taxonomic data exchanges such as the [Encyclopedia of Life](#), and to research in the *history* of science. That’s because its collections include things of interest to humanities scholars, like antiquarian books, annotated herbarium specimens, and naturalists’ field records, some 460 thousand items of which were [digitized recently](#), across multiple institutions and dating back to 1815, through a CLIR [Hidden Collections](#) grant. BHL itself has offered digitization training around the world and works with libraries and publishers not only to share open access works, but to make materials within copyright available as part of its federated corpus for research and learning.

Even when our globally dispersed inheritance of biodiversity literature is so thoughtfully drawn together (54 million pages' worth in BHL alone; 160 million taxonomic names), the resulting records are not necessarily easy for researchers to use. "[Mining Biodiversity](#)" was the theme of a productive 2015 NEH Digging into Data [grant](#), which coupled novel text-mining and visualization techniques with crowdsourcing and outreach. And projects like [PaleoDeepDive](#) and [GeoDeepDive](#) represent AI-assisted efforts to pull out so-called "dark data" from its bibliographic tar pits: those idiosyncratic features in scientific journal literature like tables and figures, that have not easily leant themselves to structured searching and the assembly of comparative datasets. Those who study the fossil record have remained, as the creators of PaleoDeepDive put it, "data limited, both in terms of the pace of discovery and description of new fossils and in terms of their ability to synthesize existing published... Many other sciences, particularly those for which physical samples and specimens are the source of data, face similar challenges."

To address issues like these, the Biodiversity Heritage Library was the beneficiary of an IMLS-funded [National Digital Stewardship Residency](#) program, which placed several NDSR residents at partnering sites, where [they worked](#) on machine learning approaches to named entity recognition and the overall computer-assisted metadata enhancement of biodiversity collections. These included projects taking up the bibliographical challenges of 18th and 19th-century field notes.

Meanwhile projects like [Digital Life](#), out of the University of Massachusetts, "aim to preserve the heritage of life on Earth through creating and sharing high-quality... 3D models of living organisms." They do this through photogrammetry, circling living creatures with their awesomely-named BeastcamTM, and converting the resulting, overlapping 2d images to highly-accurate 3d representations. And thus the field of biodiversity informatics continues to grow and pose data curation challenges of various sorts, ranging from the preservation and analysis of 3d models to large-scale environmental data generated through remote sensing, to the collection and analysis of, for instance, audio data relating to deforestation of the Brazilian rainforest. Here is a

project in which the Temb  people of Brazil are installing and maintaining what they call "[guardian devices](#)," old cell phones, hooked to solar adapters and microphones, high up in trees. These record rainforest audio and (through a machine-learning technique) increasingly accurately recognize the sounds of trucks and chainsaws up to a kilometer away, sending "a real-time alert... to the Temb  rangers, a select security force" of people from local villages who can intervene at their discretion.

The use of machine learning in monitoring contexts of various sorts is rapidly becoming the norm, and it is big business more often than community-led conservation. Microsoft has recently announced an "[AI for Earth](#)" initiative which commits \$50 million dollars in grant funds over the next 5 years for "artificial intelligence projects that support clean water, agriculture, climate, and biodiversity" and build on various APIs and shared (*ahem Microsoft*) services in the field. (Meanwhile, they invest and grow [in a more sinister mode](#).) And in the nonprofit sphere, the Mozilla Foundation just this week [announced](#) a competition to award several \$50k prizes to prototype projects illuminating the effect of artificial intelligence on society and the Web. Surely at least *some* submissions will address the ways that a healthy Internet and a healthy planet must go forevermore hand in hand.

And there continue to be efforts of a different sort, to grapple with climate change and its effects on institutional and individual cultural memory: I think here of *community organizing* work like [Project ARCC](#), through which Eira Tansey and colleagues have begun drawing together "archivists responding to climate change." And I think of [Dark Mountain](#), a collective of artists and writers "who have stopped believing the stories our civilisation tells itself. We see," they say, "that the world is entering an age of ecological collapse... and we want our cultural responses to reflect this reality rather than deny it." (If you're interested, co-founder Dougald Hine will have a note on Dark Mountain in the context of "endangered knowledge," in a special issue of the journal [KULA](#), published right here out of UVic, which I am co-editing with the brilliant Sam MacFarlane and Rachel Mattson.)

On an even smaller emotional scale, there's work like that of Kate Schapira, who is a poet and lecturer in English at Brown University. Her [*Climate Anxiety Counseling*](#) project is a Lucy-from-Peanuts-style booth she sets up in parks: "I'm scared for the effects of climate change on the world I love," Schapira writes. "Rather than try to think about, save, or mourn for the whole world, I decided to think about my city and state, and the living creatures—including other humans—who share it with me." Lately, she's been documenting the conversations she has at her counseling booth (including in a [new essay](#) on Jeff Vandermeer's *Southern Reach* trilogy), drawing local, Rhode Island organisms ([#RIorganisms](#)) on little cards to give away, and writing alternative histories of our ecological past and future. They all start with a simple phrase: "The next day..."

Deep Dreams in Digital Collections

Is everyone still with me? Now let's talk in a bare-bones way about the AI techniques that are making computationally exciting conservationist and historical projects happen—and which hold future-oriented *poetic possibilities* that explain why I set small, heartfelt and community organizing projects efforts like Project ARCC, Climate Anxiety Counseling, and Dark Mountain next to "at-scale" scientific endeavors like Digital Life and the Biodiversity Heritage Library. (And here's where, if you're a machine learning expert, you're going to feel simultaneously bored and annoyed at my reductivism for a few minutes—but when I'm finished, I hope we'll all be on the same page.)

For the past several decades, major strands of practice in cognitive and computer science worked on AI in what you might think of as a top-down or didactic mode, creating "expert systems." Think of these as hierarchical decision-making pathways, carefully designed to reflect Enlightenment principles of logic and rule-based reasoning—and then applied either to content that had been highly pre-processed to meet the machine's expectations, or on a set of thoroughly *anticipated* real-world conditions. These methods, slow to advance, were based on rather clinical and

inflexible models of (if-then) understanding. The machine “knew” what it had explicitly been told to know, and it acted in the way it was programmed.

Recent, startling leaps forward in machine learning more neatly evoke organic neural networks, or the ways perception and judgment arise organically, between the senses and the brain. And they advance not through traditional methods of computer programming, but in a kind of Darwinian evolutionary mode. Imagine a profligate process in which many little hermeneutic bots are spawned, at first almost at random, and at a shocking speed and scale: a kind of speculative computing. These critters are not carefully crafted, and they *are not smart*. But before they are permitted to breed or evolve, their ability to accurately answer a small question—for instance, to tell a typographical f from a long *f*, or a daisy from a rose—is tested against something called a truth set. This is a quantity of data, sometimes surprisingly small, which is brought into the system with the assumption that it has been accurately described, or classified, or transcribed. It’s all a little more complicated than this, but not *very* much. The bots that give correct answers to simple questions are allowed to live and propagate their accidentally good qualities of judgement into future generations—while those that miss the mark, that fail to pass their tests, are tossed out, switched off, axed.

And because we no longer *design* these little agents to understand things—we simply filter them based on their ability to pass tests—we don’t really understand, ourselves, how they work. Mostly we just understand those tests. And, generally but certainly not always, we grok the basic lineaments of the digital collection: that corpus of material the machines are being tested on. Our deep ignorance about the AI’s functioning is continually compounded, as new generations of test-passing bots build on their ancestors’ successes and train each other up to fresh challenges and ever-greater layers of complexity.

In other words, in machine learning, complex, artificial brains are self-assembling bottom-up understandings through increasingly “unsupervised” encounters with the

contents of digital repositories. Some test data are drawn from assiduously curated, historical or scientific collections. Many others come from our vast and expanding digital detritus: the traces of contemporary life—micro-transactions of various sorts—every click, every “like”—all our interactions with corporations, governments, and each other online. This is not a process that fully dispenses with human [labor](#); in contrast, it often simply obscures it, and [not to the benefit of the laborers](#) themselves, who are marginalized in many, intersecting ways. Often, we consumers are being tested, too: machine-learning algorithms are rewarded and replicated for successfully manipulating audiences—influencing us to buy a product, or watch a video, share a meme, or develop a political opinion.

For a machine to tell a daisy from a rose might require a thousand accurately-identified images of each, in a human-provided initial training set. (If you’ve never tinkered with such things, you can play with this exact idea very easily and even ultimately embed it in your phone, by following a tutorial written by [Pete Warden](#), a machine learning expert at Google. It’s called “[Tensorflow for Poets](#).”) Once trained on roses and daisies, it might only take another ten carefully labelled pictures of a sunflower or a violet for your smart phone to begin accurately identifying those flowers as well. [Ashok Popat](#), another Google research scientist who is working on handwriting and optical character recognition, recently told a group of us who had gathered at Northeastern University to discuss [multilingual and historical OCR](#) that a successful machine learning algorithm, once properly trained and working well in one language or on a set of like documents, might require only an additional *one thousand lines* of accurately transcribed text in order to retrain for a new typeface, script, or complex print or manuscript page design.

Techniques like these make evident the value of a “[collections as data](#)” approach: that is, an approach to the machine-reading readiness and the iterative, human-in-the-loop technological advancement of digital libraries for which a group of smart, IMLS-funded practitioners in the DLF community advocates. This project is led by librarians Thomas Padilla and Laurie Allen, among [others](#). The possibilities here are stunning *enough* to

realize. But more unsettling and exciting—or perhaps at least *animating*—to digital collections stewards should be the knowledge that, once set along a fruitful path, a truly successful set of machine learning algorithms can begin to produce its *own* training data to advance in understanding and pass more real-world tests. This is the generation of *completely imagined*, fictional and truly speculative collections: *manufactured* botany, or book pages—leaves that never were. It's information that the machine has dreamt up from its past encounters with real-world data, has created itself—constituted, reconstituted—in order to play the testing game in what is called a “deep learning” framework.

What's happening in these contexts is hard to see. And that opaqueness provokes a set of twinned anxieties among machine learning practitioners. These are closely aligned with calls in the broader community for greater transparency and accountability in machine learning and online representation—calls that come from folks like Frank Pasquale (author of [*The Black Box Society*](#)) and the brilliant Safiya Noble, author of the new and essential NYU Press book [*Algorithms of Oppression*](#), on how supposedly neutral search engines reinforce and reflect deep-seated structural racism:

We have to ask what is lost, who is harmed, and what should be forgotten with the embrace of artificial intelligence in decision making. It is of no collective social benefit to organize information resources on the web through processes that solidify inequality and marginalization.—Safiya Noble, Algorithms of Oppression

Noble takes a clear-eyed look at the implicit biases that permeate our digital collections and systems of cataloguing and verification, “unveiling the many ways that African American people have been contained and constrained in classification systems, from Google’s commercial search engine to library databases.” She has recently called for the regulation of search engines in the public interest, and this

summer's implementation of a new General Data Protection Regulation ([GDPR](#)) by the European Union gets us closer to that on a worldwide scale than many thought possible. In fact, Cliff Kuang [calls this](#) "a rare case in which a law has managed to leap into a future that academics and tech companies are just beginning to devote concentrated effort to understanding." Part of the GDPR demands a kind of algorithmic and machine learning accountability. Companies and other entities that hold any data on EU citizens face billions of dollars in penalties if—among other things—they do not (or cannot) share on request exactly what information they hold and how it is being processed and used.

Our "Collections as Data" colleagues are thinking along these lines as well, in the cultural heritage sphere.

*Ethical concerns are integral to collections as data. Collections as data should make a commitment to openness. At the same time, care must be taken to comply with legal requirements, cultural norms, and the values of vulnerable groups. The scale of some collections may also obfuscate what is hidden or missing in the histories **they are perceived to represent**. Cultural heritage institutions must be mindful of these absences and plan to work against their repetition. Documentation should be informed by archival principles and **emergent reproducibility practice** to ensure that users have the information they need to work with collections responsibly.*
—[Santa Barbara Statement on Collections as Data](#) (my emphases)

Now, I said I see two big machine learning anxieties that stem from and are aligned with concerns about data exploitation and structural bias. They fall under the categories, broadly speaking, of reproducibility and reproduction, or deep fakes.

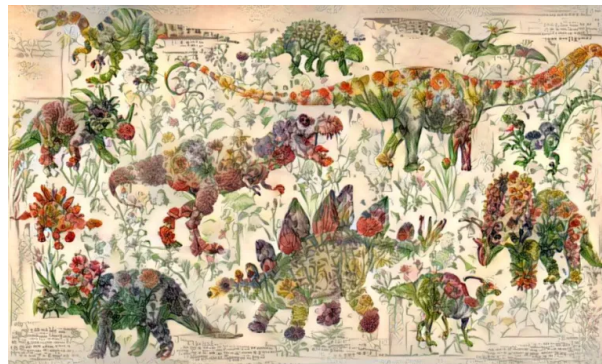
Deep fakes exploit machine learning techniques running on large audiovisual collections to produce reasonably convincing, false videos in which, for instance, a

celebrity's head might be superimposed on a porn star's body, or a politician might be made to look squarely at the camera and say things he or she would never say. A simple piece of desktop software that was released in January of this year puts the technology within almost anyone's reach, and the obvious dangers of that—not just for personal reputations but to geopolitical stability and the future of our democracies—are so great that comedian and writer/director Jordan Peele recently collaborated with BuzzFeed on a [prolonged and convincing video of Barack Obama](#), meant to raise awareness and promote greater media literacy around deep fakes. (Deep Fake Obama signs off by saying, "Stay woke, bitches.")

And then there's Lyrebird—a new, Canadian company named after an [Australian avian](#) that's a splendid and [almost uncontrollable mimic](#). A lyrebird will replicate both natural and decidedly unnatural sounds in its environment (chainsaws, camera shutters, car alarms), becoming a living broadcast system for human incursion. The [Lyrebird.ai](#) system promises to create "vocal avatars" for any English speaker with as little as one minute of recorded audio—after which, a real person's voice may be applied to any text-to-speech conversion. The company rightly touts the possibility for people with ALS or other degenerative diseases, to record and in future use copies of their own voices in machine-assisted communications devices, keeping up important emotional connections with their interlocutors even when they can no longer physically speak. Once the system advances beyond English language processing, further exciting options will open up for the preservation and transmission of endangered languages. But we well know how technologies like these are wielded by those in power against marginalized groups, and how the possibilities of fraud, harassment, and police misconduct abound when anyone can use machine learning tools and one minute of recorded sound to produce convincing vocal simulacra. Lyrebird's [ethics statement](#) boils down to, "somebody was going to do it; why not a group sincere in their intentions, like us?"

So that's an anxiety of reproduction. The other machine learning anxiety I wish to discuss is that of *reproducibility*.

Most of you will have seen or perhaps even played with “deep dream” images and generators that hit the scene in 2015—those psychedelic pictures that swept your social media feeds, in which pagodas emerged from clouds and everything that possibly *could* look like a dog did, in a kind of canine apotheosis of pareidolia. Chris Rodley’s [blossoming and burgeoning dinosaurs](#) show a further application of that same deep-dreaming technique. Rodley trained a neural network on fruits and flowers, by giving it nothing to look at but historical botanical prints—and then asked it to gaze upon some dinosaurs and show us what it saw.



Chris Rodley, “Deep Dinosaur”

What took me a while to internalize about such images is *why they exist at all*. They are in essence the byproducts of desperate attempts by developers of machine learning technologies to understand *how their own systems work*. The nature of the attempt is to run the image recognition algorithms built up independently, by all those little Darwinian bots I described—in reverse. Here’s Pete Warden on [why this matters](#): “It’s hard to explain to people who haven’t worked with machine learning,” he says, “but we’re still back in the dark ages when it comes to tracking changes and rebuilding models from scratch. It’s so bad it sometimes feels like stepping back in time to when we coded without source control.”

Neural network interpretability is becoming [a field of research](#) in its own right. Chris Olah and colleagues write that by itself, the feature visualization pictured in deep-dream images “will never give a completely satisfactory understanding,” but they see it as “one of the fundamental building blocks that, combined with additional tools, can

empower humans to understand these systems.” Cliff Kuang [says](#) of the field of “explainable AI,” or XAI, that “its goal is to make machines able to account for the things they learn, in ways that we can understand. But that goal, of course, raises the fundamental question of *whether the world a machine sees can be made to [match our own](#)*.”

This lack of understanding is something Will Knight—or, at least, his headline writer—in *MIT Technology Review*, calls “the [dark secret at the heart of AI](#).” Knight ultimately concludes that, “just as many aspects of human behavior are impossible to explain in detail, perhaps it won’t be possible for AI to explain everything it does,” and quotes Jeff Clune of the University of Wyoming’s Evolving AI Lab, who shrugs and says: “Even if [human beings] can give you a reasonable-sounding explanation [for their actions], it probably is incomplete, and the same could very well be true for AI. It might just be part of the nature of intelligence that only part of it is exposed to rational explanation. Some of it is instinctual, or subconscious, or inscrutable.”

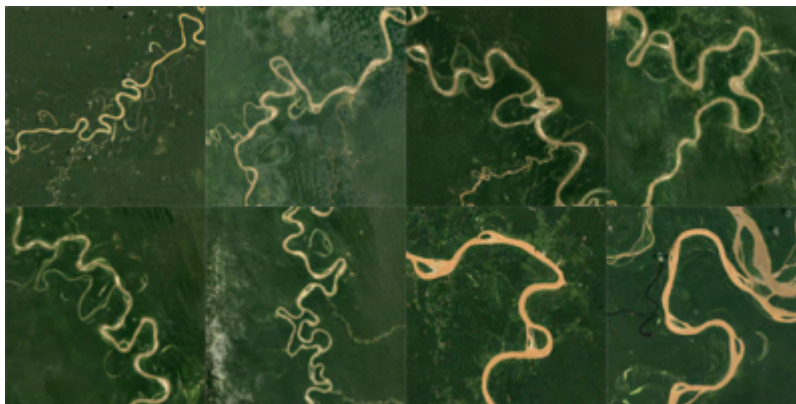
So, obviously there’s some concerning stuff here—and perhaps a reason for us, at the present juncture, to create a kind of [London Charter](#) for academic and archival work in machine learning (to promote credibility and foster professional ethics or best practices, just as the scholarly 3d visualization community did about a decade ago). Maybe next week’s DHISI [machine learning class](#) should propose a Victoria Charter? But I want to be quick to say that I also find creativity and possibility and delight in these technologies, especially vis-à-vis generative and artistic approaches to historical collections. As I attempt to pick up the pace again, let’s quickly survey a few applications of the poetic power of machine learning.

Break a Vase

[Here](#) is an MIT CSAIL project called “Videos of the Future,” in which—after watching about two years’ worth of unlabeled YouTube videos (an experiment I think my children have also conducted)—a deep learning algorithm can look at a still image and not only

predict but actually try to *create* 1-2 second clips of *the next thing that will happen*. Waves crash, a train rolls further down the tracks, golfers swing, babies do what babies do.

Abelardo Gil-Fournier is applying this technology to his artistic work on predictive landscapes, presented a couple of weeks ago as a workshop in Linz, called [*Machine Learning: An Earthology of Moving Landforms*](#). This is (I quote) “ongoing research on the image character and temporality of planetary surfaces.” As his collaborator Jussi Parikka [puts it](#), “we can experiment with the correlation of an “imaged” past (the satellite time-lapses) with a machine generated “imaged” future and test *how futures work*; how do predicted images compare against historical datasets and time-lapses and present their own ... temporal landscapes meant to run just a bit ahead of [their] time.”



Abelardo Gil-Fournier, predicted river thalwegs

[Here](#) we have Nao Tokui’s “Imaginary Soundscapes,” a “web-based sound installation, where viewers can freely walk around Google Street View and immerse themselves in an artificial soundscape [that is based on the visual qualities of real-world spaces, but has been wholly] “imagined” by... deep learning models.”

And here’s my favorite, as a lapsed Victorianist. These are Peter Leonard’s [speculative 19th-century faces](#)—a purely

creative twist on equally exciting but more straightforward work he's been doing at Yale, on projects called [Neural Neighbors](#) and [PixPlot](#). All three efforts represent machine learning analyses of about 27,000 photographs from the Meserve-Kunhardt Collection at the Beinecke. This one represents a weekend's work in training a Generative Adversarial Network. "Results seem good," Leonard [says](#). "These people have never existed."



Peter Leonard's generated portraits

When I look at rough experiments like these—or at Rodley's blossoming dinosaurs—alongside beautiful, sophisticated, and [generous](#) (but not yet, to my knowledge, machine learning-assisted) ecological design projects like Mitchell Whitelaw's biodiversity data browser, [Local Kin](#)—or the landscape concretions he describes in an important new *Open Library of Humanities* article called "[Mashups and Matters of Concern: Generative Approaches to Digital Collections](#)," I'm prompted to wonder what *more* we might do at the intersection of artificial intelligence with environmental data and natural history collections.



Alex Mordvintsev, "Remix of Ernst Haeckel illustration, generated with multi-scale style transfer"

I'd love to see, for instance, an artistic or analytical machine learning experiment using BHL collections and Scottish flower painter Patrick Syme's 1821 update to [Werner's Nomenclature of Colors](#). This book has been [recently digitized and republished](#) by the Smithsonian. It contains "the color names used by naturalists, zoologists and archaeologists through the 19th century," and it shaped Charles Darwin's formal chromatic vocabulary on the voyage of the Beagle. How might we use machine learning to identify references to these standardized colors in images and texts throughout Western library collections, and put them into conversation with indigenous color-names and perspectives on creatures living and lost?

Or perhaps we could launch an AI-assisted approach to identifying and reuniting fragmented, colonially-dispersed recordings of now-extinct birdsong with their material and immaterial or cultural traces.

Ghost Media — Kauai 'o'o Birdcall Visualization



This is [the call of the Kaua'i '??'](#)—as recorded in 1987, likely the year this Hawaiian bird went extinct. It was later further fragmented from its natural context: digitized and made available by Cornell's department of ornithology, replicated online, remixed, shared, [visualized](#), and commented on. "Events unfold themselves across centuries in random, unpredictable ways," says commissioning artist Jakob K. Steensen. "Past actions and organic occurrences become foundations for the physical realities we experience today. We live in a condition where things that happened hundreds of years ago are inherited as global extinctions, crises, and ecological catastrophes... Now, *animals* are being converted into digital, archival material at exponential rates." Can we imagine new interfaces to this kind of material—affective or scientific and pragmatic—that do not replicate violence on extinct species and on the human cultures for which they may remain in living memory?

Experiments and artistic interventions like the ones I've just zoomed through position our inherited digital heritage collections, with all their flaws and hubrises (as William James said of words and theories under a philosophy of pragmatism) "as instruments—not as answers to enigmas in which we can rest. We don't lie back upon them," [he wrote](#), "we move forward, and on occasion, *make nature over again* by their aid." It's an odd thing, I suppose, to ask the James of 1909 to speak to contemporary Afrofuturist music critic Kodwo Eshun, but I'll do it. The most startling concept I've gleaned from

Eshun's mindblowing late '90s monograph, [*More Brilliant Than the Sun: Adventures in Sonic Fiction*](#) (which I think you can also get in a nutshell, in John Akomfrah's amazing documentary of Afrofuturism, [*The Last Angel of History*](#)) is Eshun's understanding of the objects of African American cultural heritage not as things that are fixed, to be looked upon and appreciated, but as living, usable, play-able, and filled with the potential for transformation and creative re-use. In his film, Akomfrah imagines a "data thief," an archaeologist who might dig up usable code. For Eshun, the perfect example of this was a vinyl record on a turntable. By any measure, it's a recording of the past, meant for simple playback. But in the hands of a scratch artist, it becomes both an instrument and a highly accessible platform. That's the spirit in which I imagine AI might help us activate our digital collections.

Although I've described it in layman's terms (which is about all I'm capable of), I've gone into some detail today on how machine learning works because—for one thing—it constitutes, as Pete Warden [says](#), "a radical change in how we build software. Instead of writing and maintaining intricate, layered tangles of logic, the developer has to become a *teacher*, a *curator* of training data and an *analyst* of results." For Warden, this means the fundamental "replacement of traditional software with deep learning. There will be a long ramp-up as knowledge diffuses through the developer community, but in ten years," he predicts, "most software jobs won't involve programming."

Instead, they'll involve a kind of [pedagogy](#), and deep expertise not only in some problem set, area of scholarship, or subject domain but in *data curation*—in *assembling and arranging collections* of our digital cultural heritage. This is skilled archival labor, [not magic](#). If you ever needed an argument for the value and relevancy of librarianship and museum and archival studies, here it is.

Another reason I laid out the basic mechanism of machine learning is to make the point that, for all its novelty, it's not an alien process. It's in many ways deeply human and deeply connected to our bibliographical and media inheritance, and to the ways we act

as living creatures in a complex world. I'd like to transition now to a quick discussion of two things I think we hold in common with machine learning.



Learning to see: Gloomy Sunday

Memo Akten

03:02

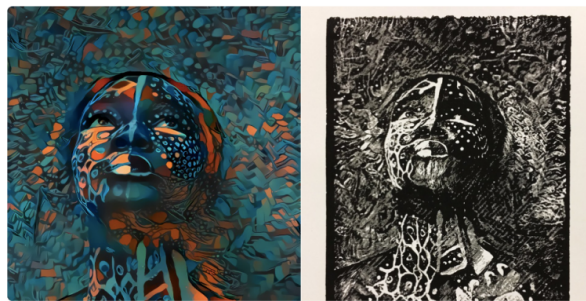
One: artificial intelligence, like us mere mortals, *only recognizes what in some way it already knows*. For instance, *this* deep neural network, now making predictions based on input from a camera, has only ever looked at ocean waves and rocks. [Watch it see the sea](#). [Later](#), Memo Akten, the artist, shows us the same system trained on fire and flowers. It's "trying to make sense of what it sees, in context of what it's seen before," he writes. Again, "it can see only [what it already knows](#), just like us."

And, two: sometimes, a neural network has to "go too far" to center itself. This is a profound insight I got from Twitter-eavesdropping on two people named Dr. Beef and The Wise Turtle. I'm [not kidding](#). Dr. Beef is a very young Stanford machine learning researcher named [Robbie Barrat](#), whose AI-generated paintings have wound up on the front page of magazines. He was recently marveling online at the way his neural network vacillated back and forth from landscapes that were dark and gloomy, or wild

and bright. Turil Cronberg [responded](#): “The organic process of learning is what I call “loopy” as it moves in a sort of corkscrew spiral, like a particle in an ocean wave. Or,” she goes on, “like a toddler learning to walk. *It has to go too far in all directions to learn how to center itself.*”

* * *

“For a long time,” [writes](#) the inimitable Rebecca Solnit, “we thought the work of climate change was imagining the future, until we realized that all our estimates were too optimistic and that the trouble was not an issue for our grandchildren but was in the present, with us, now. Even to imagine the present means summoning up the reality and the necessity of systems too vast and complex to appear before the eye. *We in the safer center had to imagine the edges.*”



Nettrice Gaskins, algorithmically generated
woodcut

I've taxed your patience enough already, but if I had more time this evening, I would tell you about digital work I see that is trying peacefully to center itself and its users (very often by dwelling in material culture and human embodiment at a moment when life seems most precarious), and about some other projects that either fail or succeed at imagining the edges from that “safer center.” These range from the amusingly-named [SkyKnit](#), a machine-learning algorithm producing knitting patterns that can be realized in yarn; to Nettrice Gaskins' [experiments](#) with woodcut and linoleum prints of deep-dreamed images; to a Snow and Ice Research Center project that attempts to learn endangered northern Arctic languages by reviewing recordings that are themselves

fixed in obsolete a/v formats; to a group of MIT researchers [who are terrifying me](#) by creating a machine learning system that picks up and understands subvocalizations—the tiny, involuntary movements and neuromuscular signals inside our faces and jaws that happen silently, when we think or read words. As with the Lyrebird technology, this tool opens up at least as many avenues of oppression and abuse as it does pathways to new and better futures.

So, once again, the scary stuff is right next to the funny stuff and the heartbreakingly beautiful stuff. It's just like... a nature documentary.

All in all, when I look around at who is best poised to take up the ideas I've shared here today in a responsible way, counter-acting *cultures of extraction and endangerment* that (just as in the natural world) often characterize the interaction of our settler colonialist institutions with the objects and lives they touch—it's committed conservationists, it's creative artists and designers, it's indigenous thinkers and holders of traditional wisdom, and it's the people working in [deeply collaborative ways](#) in living, community archives or who are informed by archival ethics [being developed in those sites](#). These are the folks I want to see advising on future-oriented approaches to machine learning in libraries and archives of the Anthropocene.

*My heart is moved by all I cannot save:
so much has been destroyed*

*I have to cast my lot with those
who age after age, perversely,*

*with no extraordinary power,
reconstitute the world.*

Poets have long shown us how to use our most ordinary powers to reconstitute the world. Because I think we're just at the beginning of a conversation and a process despite my strong sense that it's a process of assembling fragments at the end of things, I'll close not with some final pronouncement, but with Derek Walcott—some lines from his [1992 Nobel Prize lecture](#). These hopeful words, like those of Rich, helped to focus my thinking on the problem sets and fundamental fragilities I've tried to lay before you today. Walcott writes:

Break a vase, and the love that reassembles the fragments is stronger than that love which took its symmetry for granted when it was whole. The glue that fits the pieces is the sealing of its original shape...

And this is the exact process of the making of poetry, or what should be called not its "making" but its remaking, the fragmented memory, the armature that frames the god, even the rite that surrenders it to a final pyre; the god assembled cane by cane, reed by weaving reed, line by plaited line, as the artisans of Felicity would erect his holy echo.



Japanese kintsugi/kintsukuroi ("golden repair/ golden joinery") bowl. Ring the bells that still can ring.

Posted in **DESIGN, INFRASTRUCTURE**

Tagged **DIGITAL HUMANITIES, EMBODIED, LIBRARIES, SPECULATIVE COMPUTING**

◀ PREVIOUS POST

NEXT POST ▶

5 spectra for speculative
knowledge design

how the light gets in

TRAVEL/TALKS

April 23, 2021: McLeod Memorial Lecture, WUSTL, on “Cultural Memory and the Peri-Pandemic Library”

March 2020-April 2021: speaking/travel hiatus during the pandemic

February 21, 2020: featured talk, AAAD 2019: “Black Temporalities: Past, Present, and Future”

July 2019 – January 2020: speaking/travel hiatus while starting my new position at James Madison University

June 27, 2019: *Tensions of Europe 2019* keynote on machine learning & historical understanding, Luxembourg

June 19, 2019: *RMBS 2019* opening plenary on climate change & libraries/archives, Baltimore

June 2-7, 2019: teaching Rare Book School in Philadelphia: “Community Archives and Digital Cultural Memory”

March 22, 2019: *RLUK 2019* keynote on DH at the grassroots, London

THEMES

Select Category ▼

ARCHIVES

Select Month ▼

RECENT POSTS

a tribute to Stéfan Sinclair

cultural memory and the peri-pandemic library

foreword (to the past)

a pledge: self-examination and concrete action in the JMU Libraries

change us, too

from the grass roots

how the light gets in

reconstitute the world

5 spectra for speculative knowledge design

we raise our voices

iv. coda: speculative computing (2004)

inauguration day

open invitations

speculative collections

OLDIES BUT GOODIES

digital humanities in the anthropocene

asking for it

toward a new deal

resistance in the materials

too small to fail

reality bytes

lazy consensus

a skunk in the library.

why, oh why, CC-BY?

what do girls dig?

STANDARD DISCLAIMER

This site and its contents are my responsibility alone, and may not reflect the opinions of my employer, colleagues, students, children, or imaginary friends.

YOURS



Everything here is free to use under a Creative Commons Attribution 4.0 International License.



Powered by [Miniva WordPress Theme](#)