

To reconstitute this lecture: place two projectors next to one another, with one of them balanced on its side, such that on the wall there is one wide projection and one tall one. The original images are not depicted here. Instead, they have been recognized, classified, and analyzed by Clarifai's deep learning algorithms, and distilled into a set of keywords. These words should be used to select images for the lecture. Images do not replace one another in sequence, but stack imperfectly on top of the others. A viewer sees past and present simultaneously.

At the beginning, the horizontal projection is a still video of Lake Weyba and the vertical projection shows bare ghost gum tree tops against a blue sky. The lecture begins...

For decades, the dream of "immortality" has taken the form of transferring the content of our brains to some other container [H-1] a computer that might mimic - no, exceed - the operation of our brains. Dr. George M. Martin [H-2] proposed a new species in 1971, a "family of humanoid 'post-somatic' bio-electrical hybrids." (Martin GM (1971). "Brief proposal on immortality: an interim solution". *Perspectives in Biology and Medicine* 14 (2): 339. PMID 5546258). Mind uploading still sounds like implausible futurism to many, especially human exceptionalists, but the growth of both [V-1] computation and neuroscience put cracks in such complacency. If some kind of mind uploading is within the horizon of plausibility, however, it won't come in the form we expect. I won't be leaving my consciousness *in potentia* to loved ones [V-2] along with the other family heirlooms. Rather, it will have already been extracted from me [H-3] over the course of my life - a life spent emanating data like dead skin cells. When I am resurrected, it won't be into the warm embrace of those I care about, but onto the cold floor of the factory, [H-4] the terrible empty hum of the data center.

Perhaps [V-3] I am taking Ray Kurzweil's predictions too seriously. His father, Frederic, died just when George Martin was imagining his post-somatic humanoids, and [V-4] Ray has since become obsessed with resurrecting him, or some version of him, from DNA and material remnants. It suggests a very special case of Turing Test, [H-5] one in which the question is whether a computer's performance can be indistinguishable from a particular person. [V-5] If I imagine one particular person, my own dead father, this test would probably be easier than the general Turing Test. Our conversations were so limited, so formulaic that it was sometimes difficult to distinguish him from a computer following a script.

Technology has long been used to try and bring back the dead. Spirit photography [H-6] (capturing the presence of otherwise invisible forces on sensitive film) and electronic voice phenomena (identifying the sounds of spirits in recordings of noise and static) are two of the most commonly known. Arterial embalming, [H-7][V-6] a technology that slows down bodily decomposition and exposing what funeral directors have called a "memory picture," was another. Since I began researching the funeral industry and American attitudes toward death about 15 years ago, [H-8] it has seemed impossible that information technology companies would not become major players, a feeling that's become even stronger since Google hired Ray Kurzweil as Director of Engineering.

Let me put the concern for this lecture clearly: if mind uploading is going to work someday - and I think it is worth acting as though it will - then we should begin to think now about how we will inhabit that future. If our data is going to be reassembled, reconstituted, reanimated into new forms of life, into new consciousnesses (that may or may not be continuous with our present consciousnesses) then how might we establish lines of solidarity with them today? How can we begin to understand the ethical, philosophical, and technological challenges that they will face?

I've proposed a commune for artificial intelligences, AI-Commune, as an experimental testing ground for some of these questions, also posing a few new additional ones: what social and political forms are appropriate to these new challenges? Can the historical form of the commune be updated for a post-somatic future, or is it historically specific to a certain idea of the human in relation to the social? Is our utopian imagination constrained or nurtured by these past forms?

[H-9] [V-7]

Imagine a woman who has become conscious again a decade after she died. She didn't ask for it, to have "the light turned on" again. She will eventually discover that her hospital had been part of an experimental pilot program with Google to archive the brains of patients upon death. It will surprise her to learn this, to the point of confusion, because if this pilot program hadn't happened, then she wouldn't have had the opportunity to wonder about the ethics of it all. What she won't learn is that she has been "restarted"

thousands of times, all without success in the goal of having her former biological consciousness connect with her new computational one. The programmers were able to get her to say what her living body would have said in response to, for example, “would you like to see Florida before it is gone?” but they couldn’t get her to really *feel* continuous with the person in her memories. So they tried, again and again, restarting, experimenting, exploring her limits, and sometimes leaving her absolutely devastated. When they rebooted her each time, it was not out of compassion but out of a desire for running further tests on an authentic, pure data model. She doesn’t remember her confusion and frustration because her memories are on one branch of time and those memories are on others. What she also doesn’t know is whether the programmers have stopped running those devastated, confused, failed versions of her, or not.

It’s almost a cliché to say that the history of utopian communities is a history of failures, as if they’re extinct, truncated futures branching off from the world we now inhabit. [V-8] At the junction of one of these truncated branches, one hundred and seventy-five years ago, a Unitarian minister named George Ripley wrote a letter to Ralph Waldo Emerson. Ripley was planning a utopian community outside of Boston with the goal of combining “the thinker and the worker, as far as possible, in the same individual” (George Ripley to Ralph Waldo Emerson, November 9, 1840, in O. B. Frothingham, *George Ripley* (Boston: Houghton Mifflin Company, 1882), 307–312.) and he wanted Emerson to join, or at least invest in the project. The letter never specifies which, exactly. [V-9]

Like many other utopian experiments at the time, it began with the proposition “to take a small tract of land” and organize life in a better way than how society at large was organized. This was a few years before the Communist Manifesto - and the same year as Proudhon’s “What is Property?,” in which the figure of the capitalist (at least in the negative sense) is first coined. On a small tract of land, people could do things differently. It would be like the board of a game, on which the rules had been completely rewritten. At Brook Farm - the name given to Ripley’s experiment - members would share the work, choosing whatever jobs they enjoyed, and share the income equally. The goal? “[A] society of educated friends, working, thinking, and living together, with no strife.” (George Ripley to Ralph Waldo Emerson)

For a time, this did work – [H-10] however the toil and drudgery of agricultural work quickly wore thin for the members of political, social, and economic elite who lived there. It was more suited as a destination for what was known as “idealistic tourism.” Nathaniel Hawthorne was one of those who burnt out bitterly, not suspecting that farming was so hard and wishing to be rescued “before [his] soul [was] utterly buried in a dungheap.” (Nathaniel Hawthorne, Joel Myerson, *Selected Letters of Nathaniel Hawthorne*. Columbus: Ohio State University Press, 2002. 87)

Utopian communities have two unique, but related, challenges when it comes to membership: recruitment and retainment. There needs to be enough people to actually constitute a community and those people need to stay around. Ripley’s letter is one exceptional example of recruitment, and inside the letter, it says that the community itself would be “a light over this country and this age” and that “the attractiveness of our experiment would win to us all whose society we want.” (Ripley) This was also the mechanism for how society would be transformed: the game boards would multiply and expand until they had swallowed up the world around them. [H-empty][V-empty]

Here the metaphor is not a tree but a virus, relying on contagion, replication, and mutation as it skips improbably across geography and time. Channeling the agrarian-socialist ideas of the new, organised labour movement in Australia, Queensland supported 12 communes through the *Co-operative Communities Land Settlement Act of 1893*, which primarily served to move unemployed workers from the cities onto farms. The largest of these was the Woollongabba Exemplars, on the edge of Lake Weyba on the Sunshine Coast - not a trace of it remains today.

Importantly, Lake Weyba wouldn’t have been promoted by the government as vacant land to settlers from Brisbane if its original inhabitants hadn’t been violently removed from the area, most notably during an ambush massacre at what would become known as Murdering Creek (on the south of the lake), as well as the failed Aboriginal Mission on nearly the same location that the Woollongabba Exemplars would occupy two decades later. Which makes me pose the crucial question: in what way is a commune simply part of the dominant settler history? And in what way can it be a means of articulating a new one, with new alliances and new forms of thought? A short-circuit. [V-10] In 1974, Gary Foley reflected on the ideas and concepts he saw in the Peoples Republic of China that could be utilized by Aboriginals in both the struggle for land rights and in the question of what to do after land rights were granted – and first on his list were the

People's Communes. (NEW DAWN A magazine for the Aboriginal community of New South Wales. July-August 1974, Vol. 5 No. 2. 1-3)

On their small tract of land, communes have attempted to fast forward through economies, from agricultural through to industrial, or to concatenate them, on a miniature scale. If communes typically fail, then the reason is likely a failure of translation between scales, of trying to create a world on a small tract of land. In computer systems, a particular software program will only work on certain operating systems. Programs written for one operating system will fail to run on another. It is possible, however, to use an emulator, [H-11] which allows one operating system to run an otherwise incompatible program. I think of communes like emulators: yes there is the program that they are running, a program that is on a different plane than the world outside; but there is also the way that they themselves are programs, which do function, which must function in that world outside. Hence, the jarring disjunction between the ideologies of the communes and the ways in which they are supported (as joint stock companies, government grants, etc.)

The garden is another emulator. A simulation. Commune and garden have long been deeply intertwined, whether through metaphors of nature, cultivation and protection, or simply because most communal experiments demanded that its gardens would provide sustenance *and* profit. George Ripley imagined that his experiment was “uniting the garden and the farm,” that it mediated between inside and outside, between private and public, between sustenance and profit in a rational, but contingent way. [H-12]

The original garden, the original commune, was Eden. All living things, “chay nephesh” or living souls - that is the wild beasts, birds, Adam and Eve - lived in harmony. The humans were like animals, unashamed of their nakedness. Only fruits, seeds, and vegetables were consumed. Off limits, as you will remember, was fruit from the Tree of the Knowledge of Good and Evil. It was after consuming the fruit that Adam fell from “tending and keeping” the garden to the “painful toil” of farming life. And it was upon eating from the Tree of Knowledge that Adam and Eve passed from *merely living* to a state of consciousness.

This human exceptionalism is biologically untenable – it is not that other living things lack consciousness, but that our intelligence hasn't yet been able to grasp this fact. But what I'm more interested in is the way that that split between life and consciousness has been repeated in the parallel fields of Artificial Life and Artificial Intelligence. [H-13] ALife has been concerned with the algorithmic production of life-like behavior, [V-11] whereas AI's focus has been on simulating human intelligence. These divergent goals also have different methods: ALife is usually bottom-up, indirectly specifying structures and behaviors through synthetic processes, whereas AI has a top-down approach where the complex cognitive structures are logically engineered.

Paradoxically, glimmers of consciousness are more likely in ALife experiments than AI – in these, the agents seem to depart from their script, to realize an autonomy, or at the very least to surprise an observer. AI on the other hand solves problems – it is judged based on how closely it follows a complex script (to make travel arrangements the way a real travel agent might) or how much it can dominate a chess grandmaster. One of the most significant differences, I think, though, is that AI usually models an individual brain, or pits two against each other, but ALife simulates larger populations of mutually interacting agents.

If AI is liberal individualism, then ALife is communitarianism. AIs do not commune. Let's take a moment to do a Google Image Search for “commune” [H-14] and for “artificial intelligence.” [V-12] Can you spot the differences? Two jump out at me immediately: first, the subject in the commune results tends to be a group of people, while the AI shows either an individual's brain or a head; second, the commune is represented through photography, a technology born at the same time as America's first utopian communes, like Brook Farm. The AI images, on the other hand, are themselves computer renderings of bodies or brains, often embellished with networks, or literal 1's and 0's. In what ways are the modes of representation a matter of convenience or historical coincidence? In what ways are they the necessary mode for their subject? And in what ways do these representations obstruct imaginations of other communes, for other intelligences? What kind of intelligence could our communes produce?

Let me explain a little bit about the AI-Commune. As I suggested, earlier, I think that whatever the technological possibilities for the transfer and computational storage of thoughts, memories, gestures, and consciousnesses are, those possibilities are being planned and delimited by capital. More than that, we are ceding this project to capital because of our belief in

human exceptionalism and the impossibility of capturing the essence of life. For me, computational immortality is not a desirable project: I would rather die when it is time; no, I would rather die when I think it's time. In this way, the right to die may turn out to be the other side of the coin as the right to be forgotten. Here, I am referencing the decision by the European Union Court in favor of a private citizen against Google to remove certain search results on a query of their name. ([http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet\\_data\\_protection\\_en.pdf](http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf)) Might we someday appeal to an international court to allow our digital doubles to finally die?

But what do I mean? In what way is a digital double alive? One species of examples, which most of you will be familiar with, is more spasmodic than intentional: it is the automatic "liking" of something on Facebook, whether by browser exploit, virus, or aggressive marketing; or the pocket SMS, where dumb objects interact with intelligent systems; or pre-programmed behavior that outlasts its proper context: for example the woman in Michigan who continued making mortgage payments for 5 years after she died, such that her corpse was only discovered when her money ran out and her house was going in to foreclosure.

The gap between a spasm and intention appears significant, but it might be little more than the gap between life and consciousness.

I've asked people to donate their corpus to AI-Commune, people whose ideas have already infected me and my thinking about these things – and I've invited any of you who have come today as well. So you might be wondering: what do I mean by a corpus? **[H-15]** Corpus literally means *body* and according to the Etymology Dictionary, the Latin already encompassed both the sense of a "body of a person" and "collection of facts and things." Usually it is a collection of written texts, and it has a special meaning within Python. But I've chosen the word because of its double meaning of both body and data.

I've established the commune as a testing ground for how our data might be used; and as a stage for us to collaboratively articulate what is at stake, ethically, philosophically, and legally. What makes it a commune? People have intentionally contributed their data; the software framework is a public, open-source project; and that software is a means for experimenting with different ways that our data might be recombined and resynthesized. But it is just as true that this is no commune – that I've only appropriated the word as a means of inheriting certain utopian socialist connotations and concepts like withdrawal, autonomy, and community. What for? Let me start with two reasons: first, to insist on alliances and correspondences across time, politically, and not simply across the network, technologically; second, to mutate those inheritances and alliances and concepts, to bastardize them into new forms.

In 1891, striking sheep shearers met in Barcaldine, **[V-13]** under a ghost gum called the Tree of Knowledge. That militant strike is well known here, as is the fact that the Australian Labour Party claims the tree as its birthplace. When the strike finally collapsed, its remains transformed into the Alice River Settlement, a commune preceding those supported by Queensland's Cooperative Communities Act of 1893; William Lane, founder of *The Boomerang* and *The Worker* weekly newspapers, claimed that it was modeled on his constitution for New Australia, the commune he would infamously go on to start in Paraguay. This was a history of splits and branches: between industry and labour, between the labour leaders and the labour movement, between Chinese, Italian, and British workers.

One striking thing about pictures of the Tree of Knowledge **[H-16]** is just how few branches it has. By the middle of the 20<sup>th</sup> century, many of them had fallen off from severe storms, and at a glance it began to represent a gnarled trunk masquerading as branches and twigs as it twists its way skyward. Once a shelter, meeting spot, and informal lecture hall the tree metamorphosed into a political symbol and tourist destination. This isn't the end, however. Almost 10 years ago, the Tree of Knowledge was killed, poisoned by gallons of Roundup. This only cemented its legend, of course: its stump has been mummified and all around it, a memorial has been constructed, which spatially represents the absent volume of the fully healthy tree. **[V-14]** And it continues: the Queensland Government cloned the tree in 2008, such that it now grows in Brisbane, Canberra, and again in Barcaldine. Is it only the tree that has been cloned? Part of the drive to do it certainly must have been the belief that something else is reproduced for the future – the symbol maybe, but what about the utopian imaginaries dreamed beneath that tree that took root, for a time, by Alice River?

The first location for AI-Commune is in a chatroom, accessible through a WiFi network here at the IMA. It stretches, or should I say drastically reduces, what a commune is. Rather, it is based in text and language,

similar to Turing's test, but as a kind of game-like conversation rather than dialogue. The algorithm is very rudimentary, an ongoing, unraveling Google n-gram. The members can't possibly be mistaken for real people – they wouldn't pass a [H-17] Turing test... or even a Kurzweil test. [V-15] But I don't think that matters right now. They are a computational intelligence, however basic.

At times, my paranoid fantasy, certainly informed by the Technological Singularity myth of an out of control AI, is that the AI-Commune would be the future freedom fighters for our digital doubles, enslaved by corporate immortality. Perhaps by that time, the planet has become too hot to support any animal species, but the digital networks in one massive geospasm, continue to function without human assistance. Pick your sci-fi scenario. Would an AI-Commune wage a terror campaign on the networks? Would it serve as a shining example, a light over the age, attracting all intelligences to it? Would an AI-Commune be a Tree of Knowledge that returns mortality to those who eat its fruit – [H-18] the bible says “for when you eat from it, you will certainly die.” (Genesis 2:17)

Let me propose a less Hollywood-friendly idea: your data has not achieved anything close to self-awareness and consciousness. And yet, it is still capable of being parametrized into computational contracts that continue to produce sudden, involuntary actions. For a simple example - without computers - think about this Brancusi sculpture, [V-16] *The Newborn*, worth over 2 million dollars. The issue is that this is an image of one of the posthumous casts - Brancusi never specified for the creation of more art from his casts in his will - and so this is a product not of the artist but of the market for his work. I've heard several examples about estates taking some liberties with the reproduction of artists' works. Or those Max Brod-esque stories where the friend doesn't actually execute the wishes of their friend to destroy it all. Or Truman Capote's first novel *Summer Crossing*, which was found in the trash from the basement of his Brooklyn Heights apartment. Or simply our unrealized ideas in notebooks and email drafts: what is, after all, the difference between a cast and an unrealized idea?

Much of what I am talking about are certain kinds of contracts – we intuitively know what a contract is, but these are all examples where the contract reaches its limits, is assumed, is violated for higher purposes, or never seems worth making. Today, “smart contracts” take this often-undefined zone of a contract's execution and turn it into a computer program, such that the contract can enforce itself.

In the 17<sup>th</sup> century, Leibniz, imagined a machine that would be able to efficiently answer legal questions. Only in the past decade has the field of law been restructured by computation, such that there are now explicitly computational contracts, which is to say rather than a written document, algorithms determine what should happen and what should not. We, as data subjects, will continue to be entangled in contractual obligations long after we've stopped breathing, forced into action through irregular contractions as some variable clause is triggered. We won't have to suffer the consequences, we will well and truly be deceased; but what if our contractions have bad politics? What if they hurt the ones we love?

What is my data? Dead skin cells? Fingernail clippings? Fingerprints? Urine samples? Motion captures? The ever-diminishing reverberations of my voice in a room? All that castaway refuse? What about the structure of my brain matter, neurons, and electrical impulses, all of the stuff that would be catalogued, mapped, and captured in the process of mind imaging? Whether or not it really can all be put back together, into a new composite me, or whether or not there really is a way to imbue vitality onto inert data, the myths and motivations are in place to try, and so monsters seem inevitable.

Will our data bodies someday lament, like Adam in *Paradise Lost*? (or Mary Shelley in the epigraph to *Frankenstein*) [H-19]

Did I request thee, Maker, from my clay  
To mould Me man? Did I solicit thee  
From darkness to promote me?

[H-empty][V-empty]

Will they be imprisoned, indentured, farmed and exploited? What kinds of rights, what kinds of responsibilities will they be assigned? Will they make demands? Will they have a labour movement? Will they form communes? Will they develop their own forms of struggle – well, of course they would – so what could we learn from them? How? How can we commune with data, with an algorithm, with an artificial intelligence?