Artificial intelligence: Big Science, bricolage, and beyond

Stuart Watt
Turalt
Halifax, Nova Scotia, Canada

Artificial intelligence is a field with an identity crisis. Despite being over sixty years old, it is struggling to establish a coherent self-image. Is it a science, aiming to tackle the big questions about human behaviour and experience? Or is it a more practically-oriented engineering discipline, driven to build technology that can transform the world we live in?

This article will be partly history, partly analysis — and partly opinion. What I intend is that it will outline some of the shifts have happened in the way artificial intelligence works today, and how those are influenced by our social interactions within the different communities that make up artificial intelligence.

Today, much of modern artificial intelligence looks increasingly like Alvin Weinberg's "Big Science" (Weinberg, 1961). A Big Science follows the pattern of Lawrence's Berkeley National Laboratory. It is driven by large-scale funding, into the billions of dollars from both public and private sources. It is tightly integrated into the foundational economic connections of the day (in the case of BNL, the mid-20th century military-industrial complex). There are many examples in other areas of science, such as the Human Genome Project, ITER and the Joint European Torus, CERN and the Large Hadron Collider, and the US Apollo program.

In summary, the hallmarks of Big Science are: immense teams, specific but lofty goals, and massive amounts of funding.

Big Science transforms the way science is done. You can't quickly try an idea and see what happens. Even getting the point of starting an experiment may cost millions or even billions, so you need solid guidance before you take that step. The Large Hadron Collider, for example, cost nearly $5 billion and took a decade to build before it even started running experiments.

In Table 1 below, I've sketched out the characteristics that I believe distinguish a Big Science approach from a Little Science one. Of course, these are somewhat stereotyped. In practice, there is a big blurry area between them — although there's evidence of a long-term drift towards a 'Big Science culture' in many STEM fields, such that a large proportion of the work is pulled in that direction[1]. Funding calls today may be focused on grand challenges, for example, or the creation of collaborative networks.

---

[1] I'm consciously excluding umbrella-style initiatives here. For example, the EU's Horizon 2020 and the NSF National AI Research Institutes are not typical Big Science, for several reasons: they're often

"Big Science"	"Little Science"
Centralized	Distributed
Collaboration-centred	Individual-centred
Integrated into the economy	Independent from the economy
High brand value	Low brand value
> $1B funding	Modest funding
One grand challenge	Multiple, targeted questions
Planned	Responsive
Communication by press release	Communication by academic article
Applied research	Pure research
Attitude of confidence, belief	Attitude of skepticism, questioning
Consequentialist ethics	Virtue ethics
Guided by theory	Guided by experiment

Table 1. Characteristics of Big Science and Little Science

Harry Collins (2003) uses a richer typology, distinguishing "centralized big science" from "federal big science", with an overlapping mixed category showing aspects of both. It's a fascinating and remarkable history of the evolution of one example, the Laser Interferometer Gravitational-Wave Observatory (LIGO). I'd very much recommend reading it, as a beautiful illustration of how the culture and structure of the organization affects the science that is done. Big Science — however it is done — transforms the way we work together, and through that, what we do.

Is this happening with artificial intelligence today? I believe so. Many of the characteristics in Table 1 are widespread within our field, and there are organizations and initiatives that typify virtually all of them, to the point they are almost exemplary. There are quite a few initiatives that appear to the model almost exactly. Here are the ones that I first thought of:

- IBM Watson, building on IBM's competitive chess and Jeopardy, Watson was launched as an IBM business in 2014 with $1B of funding and several thousand employees. Watson is perhaps most intriguing as it 'evolved' from GOFAI to ML and integrates both. Originally intended as a universal, open-domain, natural language question answering system, it since evolved with the addition of perceptual and other techniques.
- OpenAI, founded in 2015 by, among others, Elon Musk, Sam Altman, and Peter Thiel. OpenAI's actual intent is hard to assess: originally it claimed to "democratize" AI to mitigate technological risks, but it has since pivoted to a for-profit model and closed models. OpenAI is closely integrated with both academic work (like Stanford's

---

distributed, and there isn't the unity of purpose. Research consortia formed through these initiatives may be partly shaped by these forces, but they aren't typical of them.

"foundational models") and corporate research centres in Big Tech companies like Google and Facebook.

- Alphabet's AI work is perhaps the hardest to classify, because there is an interlocking network of semi-autonomous corporations, e.g., DeepMind, and more academic groups, e.g., Google Brain. Also, Google being Google, it is hard to identify a strategy beyond "let's build stuff and see what sticks". Those differences aside, it's definitively a Google-ified version of Big Science. DeepMind was acquired by Google in 2014, and Google's AI division dates to 2017.

There's a fair case for considering Japan's "Fifth Generation Computer" project from 1982 to 1992 as another Big Science project broadly in the field of AI, albeit based on an earlier version of the field.

That covers the centralizing, community-oriented aspect, and the grand vision aspect. Over and above that, there's the cost and the funding, which is also transforming work in artificial intelligence. GPT-3, for example, cost over $10 million for a single run. AlphaFold, similarly, at cost price, would be around $7 million for training compute alone. At that scale, you need to know it's going in the right direction before you press the start button.

So, when we look at organizations like OpenAI, IBM's Watson, DeepMind, and even the likes of MIT's and Stanford's AI work, we can see a Big Science family resemblance. "Foundation models" like GPT-3 and its analogues reflect their origins in Big Science. Stanford's foundation models paper (Bommasani *et al*., 2021) has 113 authors, and has not even been formally published, yet was widely disseminated through press channels. Even OpenAI's GPT-3 and DeepMind's AlphaFold papers have over thirty authors each, and additionally acknowledge many more, including entire teams and communities. This is all very typical of Big Science.[2]

Naturally, there are problems with the Big Science approach. It doesn't always work — there have been Big Failures too (for example, Biosphere 2, and the Human Brain Project). And even when it does work, it does not always function like a *good* science. As Weinberg (one of the founders of Big Science initiatives) put it:

> "The inevitable result is the injection of a journalistic flavor into Big Science which is fundamentally in conflict with the scientific method. If the serious writings about Big Science were carefully separated from the journalistic writings, little harm would be done. But they are not so separated. Issues of scientific or technical merit tend to get argued in the popular, not the scientific, press" (Weinberg, 1961).

---

[2] These author lists are tiny by the scale of some Big Science projects. The current record is a physics paper with 5,154 authors. Thirty authors is small by Big Science standards, but it's large by historical standards in artificial intelligence and machine learning. de Solla Price (1986) discusses the impact of Big Science on publishing more than I will here, but I will note that this is an easily tested hypthesis (e.g., see Kahn, 2016).

"Foundation models" are an excellent example of this effect. In many ways, discussion of this method — and we need to be open about this, it is a methodological shift from much previous work — has had to happen *outside science* because of Bommasani *et al.*'s choice to eschew academic publishing.

Given that the defining features of true Big Science are: collaboration, grand scale, and immense funding, it seems reasonable to assert that, in these organizations at least, and those aspiring to compete with them, artificial intelligence is converging to to a Big Science model. And it is also worth noting the timeline here. The shift — if there was a shift to a Big Science version of artificial intelligence — appears to date quite specifically to around 2014, give or take.

So, how did we get there? What caused these fundamental shifts in the nature of artificial intelligence.

Kuhn's model: AI winters and paradigm shifts

Most of the people writing about how science — any science — works, inevitably start with (and often finish with) Thomas Kuhn's (1962) *The Structure of Scientific Revolutions*, a classic in the philosophy of science. It is a flawed classic, but a classic nevertheless.

Note that Kuhn's topic was *scientific revolutions*, not everyday science. He drew an explicit parallel to a (political) revolution. According to Kuhn, sciences go through a kind of lifecycle, and eventually come to a point where they stagnate, because they don't have the concepts and methods they needed to progress. When this happens, there's a kind of crisis, and the "scientific revolution" resolves that crisis by replacing the old paradigm with a new one, one which has different concepts and methods, and can continue to explore the field.

And this brings me to the second point people invariably make when writing about artificial intelligence: "AI winters". One of the more notable features of artificial intelligence as a science have been the occasional stalls, "AI winters", where funding and progress seemed to recede. To date, there have been two substantial AI winters, from around 1973 to 1980, and from 1987 to the mid 1990s.

A common account is that the AI winters were, essentially Kuhnian crisis points. Thomas Goldstein argued it: "So how did the AI winter end? It didn't! We just gave up". Others, such as Drew McDermott, have described them more as cyclical, driven by hype and disillusionment. And Melanie Mitchell, while accepting hype cycles, hedges by also conceding a fundamental shift to deep learning circa 2010.

James Lighthill, whose 1973 report arguably precipitated the first AI winter, did not describe it as cyclical, but certainly used language suggestive of a crisis: "it is unrealistic to expect highly generalised systems that can handle a large knowledge base effectively in a learning or self-organising mode to be developed in the 20th century" (Lighthill, 1973).

These accounts do not mention Kuhn explicitly, and McDermott's cyclical pattern doesn't fit Kuhn's model. But the apparently permanent shift to machine learning does, and so does Lighthill's commentary. Goldstein's narrative, in particular, is extremely Kuhnian, he even uses the phrase "seismic shift", and he was presenting to the National Science Foundation.

And Bommassani *et al*. (2021), defending “foundation models” do explicitly call them a paradigm shift (although from what, they don’t say). So I think it is fair to say that the shift from what John Haugeland (1989) called “Good Old Fashioned AI” (GOFAI)[3] to machine learning can be interpreted as a Kuhnian paradigm shift, with the AI winter as the outward manifestation of the underlying crisis.

As an aside, I’ll comment that personally, I am not convinced. I don’t think this shift was purely scientific. It is notable that both major AI winters corresponded with global economic recessions: the first with a global oil crisis and a Wall Street crash, and the second with Black Monday. So both occurred during periods of economic retrenchment, and it makes sense that investments that were not perceived as delivering enough value were cut. Economics will, inevitably, play a huge rule in the value judgements in deciding whether or not there is a crisis in the first place.

So let me tell this as a story about the transformation of artificial intelligence into a Big Science.

> **Story 1**. After about three decades of research, Good Old Fashioned AI entered a crisis over the effort involved in building systems. It was simply too expensive to manage the knowledge. This required a Kuhnian paradigm shift. That shift was towards by a machine-learning-based approach, which overcame the knowledge problem. The new, machine learning paradigm for artificial intelligence, *enabled* a Big Science version of artificial intelligence to get under way. The AI winter was a crisis and GOFAI was replaced in a *revolutionary* paradigm shift to machine learning.

But let’s look at Story 1 in detail. First off, there’s a problem with the timeline. The second AI Winter peaked around 1990, but the paradigm shift dates to, maybe if we are generous and pin it on the AlexNet paper (Krizhevsky *et al*., 2012), around 2012. A gap of *twenty two years*. That’s a long time for a revolution. And we can’t claim that machine learning wasn’t already around: it was literally driving commercial products in the 1990s. The paradigm shift (if such it was) was *not* triggered by the crisis, but by something else. So, what did happen in the 2010s that laid the foundations for it?

The second problem with Story 1 is that it can’t explain any Big Science at all. Big Science can *scale* a research tradition, but it can’t create one. All the fundamental methods and assumptions need to be in place, before it can assemble the funding and economic support systems necessary to go large. For example, you couldn’t get funding for NASA if you didn’t know it was possible to launch a rocket into orbit (NASA was re-constituted as a space agency a year after the launch of Sputnik 1). All Big Science initiatives need a proof of viability as well as enough of a social imperative before they can get started. A crisis implies an absence of any proof of viability (within the domain, at least) so we need to look for a better explanation.

---

[3] Haugeland’s term here is a nice label for the primarily symbolic approach to artificial intelligence that dominated from the 1960s to the 1990s.

Laudan's model: evolving research traditions

The problem with a Kuhnian account is that it denies the possibility of any kind of *evolution* in a science. There are essentially two states: stability and crisis. In practice, this is why more sophisticed philosophers of science, like Larry Laudan, use more nuanced models.

Barbara Von Eckardt's excellent "What is Cognitive Science?" builds on Laudan's work to map out a different model. In this model, sciences are driven by a community following one of several possible "research traditions". Because a research tradition (which is a set of methods and related assumptions) can *evolve*, sciences can progress without going through a crisis. In this sense, GOFAI and machine learning are distinct research traditions, different because while their domains may overlap, but their methods and assumptions — particularly relating to curating data and expertise — are very different.

[Note: "What is Cognitive Science?" is excellent, but a tough read, because it is extremely precise. Every word matters. I wouldn't recommend it unless you need a definitive account of the structure of the science, but if you do: it's a strong and valuable foundation.]

Von Eckardt's modified version[4] of Laudan's original structure is as follows:

1. A scientific community SC desires goal G (e.g., to explain some phenomena D)
2. The research traditions currently available to SC are RT1, RT2 ... RTn
3. The domain of a selected RT corresponds roughly to D
4. Any applied RT is not foundationally flawed (its foundational assumptions are not unsound, untrue, or without sufficient conceptual resources to explain G)

According to this model, artificial intelligence could accommodate multiple communities, using different research traditions, even for the same goal. Machine learning and GOFAI could (and did) co-exist, but, for the most part, not within the same scientific community at the same time.

This certainly happened in the interregnum between 1990 and 2010. Artificial intelligence, even GOFAI, was not dead yet. And alongside it there was, in particular, plenty of work in evolutionary computing. Early versions of word embeddings (Deerwester *et al.*'s 1988 "latent semantic analysis") were growing, though, because a scientific community could find itself in a position where an amendment to research traditions would yield improved understanding. In this case, the dominant vector-space model in information retrieval was improved by PCA-based machine learning techniques.

So, let's revise our story.

> **Story 2**. After about three decades of research, the scientific community centred around one research tradition, Good Old Fashioned AI hit a problem. The effort

---

[4] Von Eckardt's primary modification from Laudan's is to drop one factor, whether or not a chosen tradition is required to have a higher rate of progress than the others. This makes sense, and also explains why empirical and theoretical communities could exist, side by side.

> involved in building systems was becoming uneconomic. It was simply too expensive to manage the knowledge. One of its foundational assumptions, that it was economically viable to capture and use knowledge to build applications, was found to be false. Other research traditions, not incorporating this assumption, were less affected. Accordingly, the GOFAI community evolved, in various ways. One new community embraced the problem, and branched into *knowledge management*[5], including early *recommender systems*. Another related branch focused in improving methods, developing methodologies like CommonKADS. Over time, machine learning techniques, particularly those that assisted knowledge acquisition, became an established part of those evolved research traditions, and GOFAI simply faded out.

There are points when abandoning a technique is *rational*, i.e., when — given the available technology — other techniques perform better. i.e., the exact same technique in 1973 might be less useful than in 2013, simply because in 1973 it might be prohibitively expensive. Convolutional neural networks, introduced in 1980, the era of the Intel 8088, are bound to have a different outcome in the era of NVIDIA GPUs. So, a shift in the cost of knowledge could have far-reaching effects on the direction of research.

Sadly, this doesn't explain everything either. It does account for the early evolution of knowledge management and machine learning/information retrieval technology, which grew dramatically after 1990. But it doesn't really explain the transformations around 2010, as well as a more Kuhnian paradigm shift would. And it does leave us with some rather intriguing questions about the apparent death of, e.g., genetic algorithms, which had a very distinct peak around 1998 and *also* dropped off the grid by 2010, despite no obvious flaws in any of the fundamental assumptions.

However, this is where Von Eckardt's tweak to Laudan's model is important. She dropped the requirement that a community select the research tradition that optimizes progress. And she's right. The people in a scientific community have different goals, like securing funding, promotion, and publications. In a Big Science world, these matter more than in Little Science. So the selection of research tradition optimizes personal career prospects as well as the scientific goal G.

Put simply: if it's easier to get grants and publications in deep learning than in genetic algorithms, people are going to switch. In other words, factors like cost, efficiency, and how easy it is to get published or promoted, absolutely structure the selection of research traditions. It is not only about the science.

---

[5] The timing here is significant. The dramatic growth in knowledge management started *exactly* as the 'AI winter' was hitting its height, between 1990 and 1993. And the central assumption of knowledge management was that *knowledge is expensive*. The change in foundational assumptions coincides perfectly with the fall in expert systems and the rise in knowledge management. Also, many people switched from one field to the other (I was one of them!) This was, at least partly, a change in *positioning*, not a change in direction. However, it also prepared the ground for increased use of machine learning.

We can see more clues in what happened in the aftermath of Lighthill's (1973) commentary. Lighthill identified three categories of work in artificial intelligence at the time: automation, cognition, and a big mushy area between the two that included robotics. He found good evidence of positive results in the first two categories. However, he held that the foundation of artificial intelligence depended on the assumption that these were a unified continuum, and, unfortunately, little of what was happening in the mushy middle was perceived as delivering value, at least at the time. That was where the axe fell.

So what happened was a branch point, if you like. Many existing areas continued, perhaps more independently of any AI umbrella. This included automation, and fields like information retrieval. Cognitive science also branched out as a newly unified and distinct field, thanks to Longuet-Higgins's commentary on Lighthill, and with significantly closer links with the social sciences (especially psychology and economics). What does this branching mean for artificial intelligence?

Mulkay's model: branching research traditions

When we look at a field like artificial intelligence, it is easy to imagine — thanks to the Big Science framing — that it is a coherent and relatively unified thing, where our similarities outweigh our differences.

This is a long way from the truth.

Michael Mulkay (in the 1975, "Three Kinds of Scientific Development") extends the Kuhnian model (which he calls the "model of closure") by showing that despite the shared identities and methods and assumptions — on the ground, a science is made up of many smaller communities. Mulkay argues that research traditions (which, remember, include methods as well as assumptions) evolve when people migrate between communities, taking ideas and methods with them as they go. And, migrations don't only happen *within* a field, but more importantly, *between* fields.

For example, in artificial intelligence, some of the fields include:

Image classification; image segmentation; speech segmentation; image generation; text generation; speech synthesis; planning; optimization algorithms; unsupervised learning; reinforcement learning; computational biology; word sense disambiguation; conversational AI; generative models; adversarial networks; recurrent networks; time series analysis; causal inference; robotics; decision making; regression; dataset construction; safety; privacy; interpretability, and so on.

At a typical conference, you might see these manifest as "tracks". These are the scientific communities that make the field work, and they typically comprise maybe a couple of hundred people. By and large, people prefer to stick with one field, and will personally know most of the others within that field. But they generally won't work across all: few people will work full-time on both, for example, speech synthesis and computational biology, or image generation and causal inference. However, people may move between fields from time to time, perhaps because they change job, or meet a new collaborator.

There are many parallels between the Laudan/Von Eckardt model and Mulkay's (not a huge surprise, Laudan and Von Eckardt are philosophers and Mulkay's from sociology, so they might well observe similarities). But Mulkay is clearer about the level of granularity, and about the mechanics of the evolution of research traditions. *Concepts and methods* move between research traditions because *people* move between research traditions.

Let's see if we can use this to improve our explanation.

> **Story 3**. After about three decades of research, the scientific community centred around one research tradition, GOFAI expert systems hit a problem. The effort involved in building systems was becoming uneconomic. It was simply too expensive to manage the knowledge. One of its foundational assumptions, that it was economically viable to capture and use knowledge to build applications, was found to be false. Other research traditions, not incorporating this assumption, such as the neural network community, were less affected. The GOFAI expert system communitybranched into knowledge management, early recommender systems, and methodologies, repositioning itself to maintain funding. Time passed. One day, a neural network researcher saw a departmental seminar by a computer vision researcher, who was using GPUs to greatly improve performance. They adapted the idea to neural networks, and it showed promise, as they gradually scaled to bigger datasets. By the early 2010s, GPUs had provem to make the high compute costs of machine learning much more affordable. And by the mid 2010s, there was proof of viability via AlexNet, proof of need in recommender system usage in Amazon etc., and a step change in compute economics — enough to create conditions for a Big Science scale investment.

The first observation about this story is that the "crisis" and the "paradigm shift" are decoupled, allowing us to account for the twenty-year gap between the two. In fact, the first half of the story is identical. The difference is the account of how deep learning got its traction: it was not a reaction to a crisis, but an (initially small-scale) exploration of a possibly-valuable method.

Secondly, this account doesn't work at the 'artificial intelligence' level, but eventually reaches it through a bottom-up spread of buy-in, as the new methods spread virally between communities.

This story is not randomly made up. If we think of AlexNet again in this light, its true innovation was the re-purposing of GPUs, but that work was already under way at another lab *in the same institution*.

Now I can't prove this, but I will bet the grand sum of $10 (Canadian) that the actual innovations came from the OpenVIDIA project on computer vision, which, like Hinton's group, was also based at the University of Toronto, but in the EyeTap Personal Imaging Lab. OpenVIDIA, developed from 2004 to 2006, was leveraging GPUs for image processing. (The EyeTap, developed by Steve Mann, was a precursor to Google Glass.) Anyway, OpenVIDIA embraced CUDA as it emerged in 2007. Cudamat (Mnih, 2009), also at Toronto, then integrated Python with CUDA, and Alex Krizhevsky was using CUDA in 2009, and Ilya Sutskever by 2010. And after all that, applying the same methods to ImageNet is not a big leap, and certainly not a revolutionary reaction to a crisis in GOFAI. It is appropriation of a useful method, through what is very likely personal contact at the same university department, and some reciprocal exchanges between OpenVIDIA and NVIDIA.

Personally, I love the insights coming from the likes of Collins and Mulkay, looking not at how science *ought to* work, but at how it *does* work, in practice, in the ground. Harry Collins provided some of the most thoughtful and valuable commentaries on GOFAI, in 1990's "Artificial Experts", and again on language in AI in 2018's "Artifictional Intelligence" (much of which is directly relevant to large language model work). Both are well worth a read. I truly hope that constructive criticisms like these continue from the social sciences as well as within the field.

The alternative: bricolage in science

Not all science is Big Science. Even Weinberg, arguably its biggest advocate, was explicit that "We must make Big Science flourish without, at the same time, allowing it to trample Little Science" (Weinberg, 1961).

However, I want to add to Table 1 above, one more distinction between Big Science and Little Science: where Big Science uses engineering, Little Science uses *"bricolage"*.

Bricolage (a loan-word from French, roughly equivalent to the English "DIY" or "tinkering") was introduced as a concept by Claude Lévi-Strauss in *The Savage Mind*. He uses it to describe how conceptual structures are put together, piece by piece.

Bricolage, put simply, means playing around, trying ideas, and testing them. It's an effective problem-solving strategy, although more so in the absence of guiding knowledge. Engineering, by contrast, is thinks about goals, and means to an end, i.e., using knowledge first.

Unhelpfully for us, Lévi-Strauss contrasts bricolage with "science" more than "engineering". Essentially, his distinction is that the bricoleur pieces together conceptual structures from (often second-hand) observations, bottom-up, where the scientist interprets observations from conceptual structures, top-down. Although, as Lévi-Strauss himself put it, "both approaches are equally valid" (Lévi-Strauss, 1962, p22). What is confusing for us is that Lévi-Strauss's use of "science" is more of an ideal than a reality. If we look at Goldstein's distinction between "science" and "principled machine learning", it exactly matches Lévi-Strauss's, except that science has 'reversed its polarity' — Goldstein's science, and ours, especially Little Science, happen substantially through bricolage — and therefore *align with* science, rather than opposing it in Lévi-Strauss's original usage. Instead, I'd contrast bricolage more with engineering, where there is a more 'principled' construction — after all, contrasting do-it-yourself with engineering does make more sense.

In many ways, it is better to use Seymour Papert's 1993 adaptation of bricolage as a method for building mental constructs, i.e., *learning*. Learning happens through bricolage, and whether it is creating new concepts and methods in science, or new concepts and methods in a child's mind, the process is virtually the same. For example, if we revisit the later parts of Story 3 through the lens of bricolage, the "do it yourself" assembly and re-purposing of tools and ideas is extremely clear.

The problem is: bricolage doesn't work well in Big Science. It can't. When you are building something the scale of GPT-3, the Human Genome Project, the LHC, or the Manhattan Project, experiments have to be few and far between. The scale, and the cost, simply make tinkering unacceptable or uneconomic, and usually both.

Is Big Science AI the child of capitalism?

It could be argued that that Big Science is what happens when capitalism gets engaged with science, and given my examples (OpenAI, etc.) this is an point that deserves serious consideration. But is it actually true? Is Big Science really the child of capitalism?

Much of our world has been transformed by Big Sciences. They can be extremely valuable in all sorts of different ways, even incidentally. Many of the classic examples of Big Science include NASA, CERN, the Human Genome Project, have all generated results that are transforming our lives. And none of them would have been possible without very substantial investment, in one form or another.

But are they capitalist? Not necessarily, not intrinsically. Some of the older Big Science projects, e.g., NASA's Apollo programme, CERN, were very much not-capitalist. And Big Science was also a big deal in the Soviet Union. The space race was driven more by international competition and the Cold War than by capitalism. The Human Genome Project, too, was not very obviously capitalist.[6]

However, Big Science does reflect a fundamentally different science. The way science is done today, at all levels, can be traced to the structural changes of Big Science. For example, Weinberg (its biggest advocate), argued in 1961 that two necessary changes were:

> "First, a great expansion in the use of short-tenure, postdoctoral fellows at the big laboratories, and second, the establishment of independent graduate schools of technology in close proximity to the big laboratories, and with some interlocking staff." (Weinberg, 1961)

These changes are both now intrinsic to the way science is done today, and both have rewritten the career system radically for those involved.[7] Not in a particularly good way,

---

[6] I admit that the artificial intelligence manifestation of Big Science does *seem* substantially more capitalist. But, recall that all Big Science initiatives do have tight economic integration, and always have had. And that capitalism does have a different attitude to innovation. I will come back to this point shortly, in the light of Ulrich Beck's model of the risk society.

[7] Especially so in North America. Big Science is (and always has been) as much a political project as a scientific one. International competition drives it at least as much as national pride in science. In fact, there is always an international dimension to any true Big Science, although it may be either competitive or collaborative, and sometimes both.

either. Big Science, being structured around organizations, often converges on a pyramidal, hierarchical structure. And like an iceberg, only the tip may be visible.

So how do we explain these global changes? Perhaps the tilt towards Big Science, what we might call "bigsciencification", is an effect of "late modernity". Science is directly linked to risk, and late modernity's management of risk (what Ulrich Beck calls the "risk society"). The changes in science, its democratization and secularization, its diffusion into wider society, are effects of its transformation under late modernity.

Beck is worth quoting in full:

> "Science, having lost reality, faces the threat that others will dictate to it what truth is *supposed* to be. That is not only the case with the flourishing 'court science', by way of direct influence. The approximate nature, the indecisiveness, the accessibility to decision-making of the results make this possible. Selection criteria that escape scientific scrutiny achieve a new and perhaps decisive meaning in the hypercomplexity that must be mastered in any case. These include the compatibilities of basic political views, the interests of sponsors, the anticipation of political implications; in short, political acceptance" (Beck, 1992, p167-168, original emphasis)

Much of what we see in artificial intelligence matches this pattern. Gone is the monopoly on knowledge within universities. Wider integration into sponsorship and governments replaces it. But there is a cost. As Beck puts it:

> "This is a development of great ambivalence. It contains the opportunity to emancipate social practice *from* science *through* science; on the other hand it *immunizes* socially prevailing ideologies and interested standpoints against enlightened scientific claims, and throws the door open to a feudalization of scientific knowledge practice through economic and political interests and 'new dogmas'" (Beck, 1992, p157, original emphasis)

Beck's thesis also accounts for the transformation of the Big Science approach, from orthogonal to capital in the first half of the 20th Century, to aligned with it today. So, I'd argue no, this is not simply capitalism. If it is anything, it is late modernity, which is *also* transforming capitalism (not in a good way). The causes run deeper. It's not simply a case of modern artificial intelligence, to borrow from John Wyndham, being "too contaminated by capital to keep afloat".

But that is not to belittle the problem. There *is* a problem, a big one. Science is searching for new structures, and Big Science seems to be winning at the moment. And we can see the signs of feudalization everywhere, from the armies of post-doctoral researchers, to the hyper-wealthy owners of private research monopolies.

Without serious work on the part of scientists, social scientists, governments, and industry, this will push out the bricoleurs who will create the next generation of innovations and discoveries.

Big Science and ethics

Ethics is different in the context of a Big Science. As I suggested in Table 1 above, one of the differences is that, in Big Science, ethics is no longer an individual matter[8]. Generally, it seems that ethical decision-making tends to be more consequentialist in a Big Science. This has permitted relatively heinous ethical actions, including, for example, the abuse of Henrietta Lacks' DNA, and the creation of the atom bomb.

This pattern is also visible in artificial intelligence work. Self-driving cars are indubitably a Big Science project[9], and the consequentialist *narratives* are clear. (Here, I'm talking about *researcher* ethics, not the ethics that they look to implement, i.e., what shapes the "self-driving project", not how a self-driving car should operate.)

As Hardin argues, Big Science *needs* external formal ethical regulatory systems and processes in a way that Little Science may not. This is inevitable given the conflicts between organizational and institutional interests and scientific goals. Sadly, for the most part, regulation in artificial intelligence is not heading an ideal direction; at present it seems to be grounded primarily in governmental rules. We need *processes*. And we need mechanisms that do *not* result in people getting fired for alerting the community to ethical concerns.

To return to Von Eckhart's model, a Big Science's goals will invariably be partly non-scientific, driven by the need for perpetuation of the institution. For example, OpenAI's goal is not (only) pursuit of truth, or building good technology, as much as perpetuation of OpenAI itself.

The future of AI in a Big Science world

We are, whether we like it or not, in a world which has tilted towards a Big Science model, across all sciences. But this is particularly and especially true in artificial intelligence. This is a problem. On this point, Goldstein is right: some rebalancing is needed to promote bricolage/experimentation/Little Science, to keep the flow of innovation going in a world which is flooded with Big Science tech, like foundation models. Without this rebalancing,

---

[8] Careful observers will have noticed that I made no defence of the ethical differences between Big Science and Little Science at that point. I still haven't — although I think it is a fair observation. If pressed, I'd argue that this follows from the institutional/communal structure. Individual ethical frameworks are inevitably less significant when dealing with a group. And there are some powerful dynamics, like "collective narcissism" (Golec de Zavala *et al*. 2018), which can trigger extreme hostility to criticism. I think communal narcissism is one possible reasonable explanation for the utter fiasco of Google's treatment of its own AI ethics people, which was triggered initially by fair criticism from within (Bender *et al*., 2021). However, the point is that as an institutional/communal enterprise, any Big Science introduces psychosocial forces which (a) may not be present in Little Science, and (b) may not be conducive to *good science*.

[9] I've not touched on AGI, which is an even bigger Big Science project in its grandiosity. However, I'm not certain it has risen to the Big Science threshold yet. There is neither the massive funding nor the unity of purpose or investment to drive it.

Big Science will stagnate, and we cannot sustain the expectations. In fact, given the sheer amount of high-brand-value, innovation-by-press-release, we might not be able to anyway.

But I believe the mechanics of the process are different. AI winters — such as they exist — are not Kuhnian in origin. Instead, the dynamics of Big Science interrupts the ability of communities to inter-connect. The stagnation is more the consequence of the Big Science institutional pattern 'freezing' existing research traditions and inhibiting the creation and evolution of new ones that compete with them.

We need to accept that multiple scientific communities with multiple research traditions is a Good Thing — this is why, for me at least Story 3 has a plausibility that Story 1 and Story 2 do not. We need to fund work other than deep learning, other than neural networks, and even — shock! — other than machine learning — if we are to build a strong and vibrant field, robust from hype cycles.

But this is about more than funding. The challenge is not, despite calls from those like Weinberg and Lauer (2014), to "achieve the right mix". It is: how do we preserve Little Science in a dominant culture of Big Science, i.e., when the values of Big Science are considered "normal"?

I can think of four good things we can do — I am sure there are plenty more, but this is where I would begin.

Thing One: empower and support the bricoleurs. The easiest way to re-balance Big Science and Little Science would be to give some of the resources to Little Science. There are many ways to do this. One would be as simple as a universal basic income. Remember, science has become more democratized, citizen science is a thing now. Another is to build a better career system (in fact, a real career system) for post-doctoral researchers, who have been true victims of the interaction between late modernity and Big Science's fundamental career changes. More of a challenge is how to truly open and democratize the networks, but, for example, opening up conferences and workshops would be a start. Make it easy, and cheap, not only for people to *use* artificial intelligence (after all, the software is essentially free now) but to *participate in it*.

Thing Two: bring forward a humanized, Slow Science, Slow AI. Yoshia Bengio has explicitly pointed to "Slow Science" (which very much matches the branching, bricolage-oriented, Little Science style). He does this as a reaction to the Big Science problems in machine learning.

Thing Three: reframe machine learning as a transdiscipline. Transdisciplines are a special kind of science. Michael Scriven describes them as: "a discipline that has standalone status as a discipline and is also used as an methodological or analytical tool in several other disciplines" (Scriven, 2008). His examples include: statistics, logic, design, and communication, with possible evidence for ethics, computer science, and information science (and, therefore, information retrieval). There is a lot to be said for artificial intelligence, and *especially* machine learning, as transdisciplines. It transforms the relationship to related sciences and fields, such as engineering, linguistics, medicine, and psychology: strengthening connections between them. Machine learning needs to learn this from statistics — what I'd call the queen of transdisciplines. Machine learning folks need to build tight collaborations with people in other fields, while still preserving their identity and innovation as a field in their own right. Until they do become a transdiscipline, this XKCD will continue to apply.

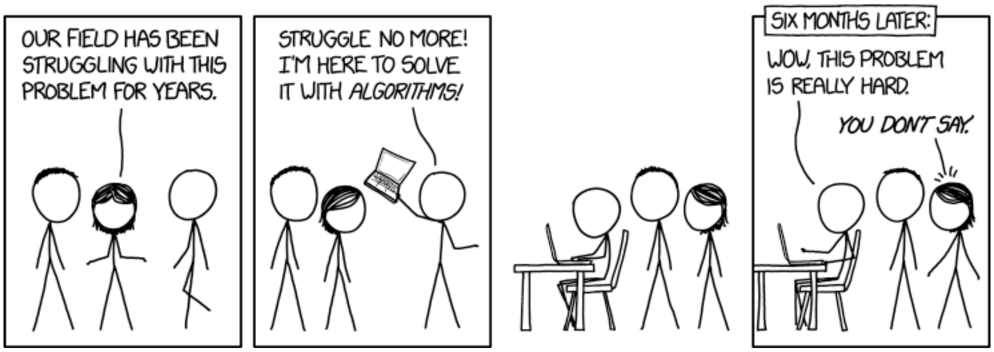

Thing Four: stop worrying about AI winters. We need to stop thinking about machine learning as a Kuhnian paradigm shift from GOFAI. It's a set of methods — a very powerful set of methods — that were unlocked by technical innovations. I am glad they were unlocked, because I like them. But a Kuhnian model is unhelpful for several reasons. First, it explicitly regards all past artificial intelligence work as, essentially, junk (see, for example, the way Goldstein put it). This is both false and socially corrosive. It is good to be conscious of hype cycles, and we need to be far more critical of the way Big Science feeds them. When we look in more detail, as in Story 3, we see smaller communities exchanging ideas and methods in a more dynamic way.

I am sure Big Science is here to stay. It is too deeply intertwined with modern scientific practices to fall away any time soon. But, I hope you will think about it, and its effects on the way we work and the way we interact. Big Science can make real things which would otherwise be impossible. Large language models — for all their inherent problems (Bender *et al*., 2021) — are vital to study; we need to learn from them, even if we must remain critical of their many biases and shortcomings. However, as Weinberg (1961) said, even back then, we do need to distinguish scientific from popular commentary, while remembering both have a part to play. The thread from the Big Science approach is that it removes entire topics from scientific discourse and plays them out in public alone.

So, I hope you will remember that, in practice, Big Science is not enough. We need the bricoleurs, we need the Little Scientists, because these are the ones providing the guidance, the reflections, and the pieces for the next groundbreaking changes in our understanding.

Long live the bricoleurs, may their tinkering be forever rewarding.

Afterword*: Thank you for reading this far. I'm planning on some more quantitative analysis of work in the field, maybe even some qualitative too.*

If you choose to support me through @buymeacoffee, I'd be both grateful and motivated.

***Disclosure**: I've worked on several Big Science projects, including the International Cancer Genome Project and AACR Project Genie.*

References

Baneke, D. (2020). Let's not talk about science: the normalization of Big Science and the moral economy of modern astronomy. *Science Technology and Human Values*, **45**(1), 164–194.

Beck, U. (1992). *Risk Society: towards a new modernity*. London: SAGE Publications.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: can language models be too big? In the proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (Vol. 1). ACM.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., *et al*. (2021). On the opportunities and risks of foundation models. arXiv preprint: http://arxiv.org/abs/2108.07258.

Collins, H. M. (1990). *Artificial experts*. MIT Press.

Collins, H. M. (2003). LIGO becomes big science. *Historical Studies in the Physical and Biological Sciences*, **33**(2), 261-297.

Coppola, N. W., & Elliot, N. (2005). Big Science or bricolage: an alternative model for research in technical communication. *IEEE Transactions on Professional Communication*, **48**(3), 261–268.

de Solla Price, D. J. (1986). *Little science, big science... and beyond*. New York: Columbia University Press.

Deerwester, S., Dumais, S., Landauer, T., Furnas, G., & Beck, L. (1988). Improving information-retrieval with latent semantic indexing. In the proceedings of the ASIS annual meeting (Vol. 25, pp. 36–40).

Esparza, J., & Yamada, T. (2007). The discovery value of "Big Science". *Journal of Experimental Medicine*, **204**(4), 701–704.

Goldstein, T. (2022). Deep learning needs more science! Or... can deep learning end the AI winter? Presented at NSF CIF Town Hall, January 10 2022. See also: his Twitter threads summary.

Haugeland, J. (1989). *Artificial intelligence: the very idea*. MIT Press.

Kahn, M. J. (2016). Big Science, co-publication and collaboration: getting to the core. In the proceedings of the 21st International Conference on Science and Technology Indicators. Valencia, Spain, 653-660.

Kuhn, T. S. (1970). *The structure of scientific revolutions* (2nd ed.) University of Chicago Press.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, **25**.

Lauer, M. S. (2014). Personal reflections on Big Science, Small Science, or the right mix. *Circulation Research*, **114**(7), 1080–1082. Lévi-Strauss, C. (1984). *The savage mind*. Man and World (Vol. 17).

Lighthill, J. (1973). Artificial Intelligence: a general survey.

McDermott, D., Waldrop, M. M., Chandrasekaran, B., McDermott, J., & Schank, R. (1985). The dark ages of AI: a panel discussion at AAAI-84. *AI Magazine*, **6**(3), 122–122.

Mitchell, M. (2021). Why AI is harder than we think. Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '21), p3.

Mnih, V. (2009). Cudamat: a CUDA-based matrix class for python. Department of Computer Science, University of Toronto, UTML TR 2009–004.

Mulkay, M. J. (1975). Three models of scientific development. *The Sociological Review*, **23**(3), 509–526.

Papert, S. (1993). Rethinking School In The Age Of The Computer. In *The Children's Machine* (pp. 138–156).

Scriven, M. (2008). The concept of a transdiscipline: and of evaluation as a transdiscipline. *Journal of MultiDisciplinary Evaluation*, **5**(10), 65–66.

Von Eckardt, B. (1992). *What is cognitive science?*, MIT Press.

Weinberg, A. (1961). Impact of large-scale science on the United States. *Science*, **134**(3473), 161–164.