

Re:Generating the Tools of Knowledge

Alan F. Blackwell
University of Cambridge

*Draft paper for presentation at the fourth workshop on
Science in the Forest, Science in the Past in June 2024*

Abstract

Contemporary Western science has configured knowledge as the production of disembodied text, for years in the form of academic publications, and more recently in the enterprise of artificial intelligence. Information technologies encode knowledge as a commodity to be generated and exchanged. Recent advances have largely met the challenge posed by Alan Turing, in which the re:presentation of intellectual labour via a sufficiently mechanised system can be hard to distinguish from online social interaction. Having framed the generation of knowledge as its re:generation, we urgently need to take stock of what remains distinctively human. Following many generations in pursuit of the ideal “rational animal”, future re:generations need to become embodied creative agents, interacting with knowledge tools, illustrated here with an auto-ethnographic reflection on the music of the live coding drone metal band Chainsaw Trousers.

Introduction

Science is concerned with the accumulation of knowledge, and with its generation. I should hasten to prefix this observation with “it seems to me that...”, since I am neither a scientist, an anthropologist, or a historian or philosopher of science, but an engineer, whose first publication appeared in a journal with the metaphysically provocative title *Knowledge-Based Systems*. In the years since I entered this field, we have heard much of “knowledge workers” in the “knowledge economy”, where public investment in STEM (Science, Technology, Engineering and Mathematics) is expected to contribute to national wealth through the production of this valuable commodity, whether knowledge

productivity is achieved by the generation of new knowledge or by acquisition of knowledge being accumulated online or generated elsewhere¹.

Who wouldn't want to be wealthier, and to have more knowledge? Government policy and investment is preoccupied with increase of wealth - with economic growth - and with the measures of productivity that are expected to deliver this. In the knowledge economy, the generation of knowledge seems inarguably a public good delivered through science. It would be absurd to argue for *less* knowledge, apart perhaps from the perspective of Kurt Vonnegut's satirical novel *Galapagos* (1985), set a million years in the future when our aquatic descendants have abandoned the big brains that turn out to have been an obstacle to streamlined ocean life, and after an AI-like device called the Mandarax has so angered the very last ship's captain, with the uselessness of all its knowledge, that he throws it into the ocean.

From the perspective of economic planners, increase in wealth is associated with productivity - the rate of production of some measurable substance or commodity resulting from human labour. A typical statement of this assumption is the following rather routine quote by think tank economist James Smith, research director at the Resolution Foundation, in a BBC radio interview:

"We're working harder, we are producing more [...] but the growth in how much more we're producing has slowed down markedly, and that is really the key driver of long-term improvements in how wealthy we are as a country." BBC Radio 4 *More or Less*, 21 Feb 2024

Smith doesn't mention *what* is being produced, only how much². For an economist, the details of one commodity rather than another may be secondary to his fundamental interest in wealth itself. Nevertheless, we might suspect that, from this economic perspective, the economic value of any knowledge worker (including the author of this paper, and perhaps all of you who are reading it) can be measured according to *how much* knowledge they are accumulating or generating, whether or not some kinds of

¹ A typical sentiment, in a letter to the Guardian following debate about the financial sustainability of UK universities, is expressed by self-described 91-year-old emeritus professor Michael Bassey as follows: "Universities have three equal purposes. First, to accumulate, store and disseminate useful knowledge; second, to engage in research to increase useful knowledge; and third, to help students acquire useful knowledge beyond what they have learned elsewhere."

² In the words of songwriter Rosie Tucker "The metadata proves you're a real thing [...] If it can't be counted, does it count?" from Big Fish / No Fun, on the 2024 album UTOPIA NOW!

knowledge³ were suspected of having no intrinsic value. Furthermore, when prioritising future investment in science (as in any other investment), the *amount* of knowledge, as this economist states so clearly, is less significant than the *rate of growth* in that amount. Setting aside the quibbles of satirists like Vonnegut, the best science must surely produce more knowledge, producing it faster and in greater quantities.

The Mandarax in Galapagos was by no means the first time Vonnegut had satirised computers and the accumulation of knowledge in the digital economy. Indeed, his very first novel, *Player Piano* from 1952, described “America in the Coming Age of Electronics,” where the Player Piano of the title is a punch card-driven automaton replacing human labour. Although the electronic giant brains causing such concern in that decade were too large to be carried on a boat, let alone hurled out of one, the problematics of bigger brains were all too clear, whether machines or the people being replaced by machines through the economic logic of productivity growth. Even more worrying, 70 years after *Player Piano*, the environmental costs of today’s largest AI research projects (Crawford 2021, Bender 2021), suggest that these massive brain simulators may actually be hastening the evolutionary replacement of homo sapiens as imagined by Vonnegut.

Vonnegut’s novels are not, of course, science in the past, but science fiction *from* the past. And we should be cautious about heralding the science fiction writer in retrospect as prophet or inventor. Nevertheless, the Western adoption of routine computerisation throughout the 1980s offered an observational context within which to assess the concerns that Vonnegut developed 30 years before that, and we have now reached a point where AI large language models (LLMs) such as ChatGPT from OpenAI are so close in capability to the Mandarax described by Vonnegut, that we can ask seriously what such developments mean for our evolutionary trajectory as a species, the size of our human and digital brains, and the utility of the knowledge we are accumulating and generating.

Re:generating knowledge

At the SFSP III workshop, my colleagues and I presented work that was very specifically concerned with the accumulation of knowledge, studying data science in

³ Perhaps even theories of economics?

the forests of Africa and Latin America as a potentially extractive and neo-colonialist enterprise undertaken by Western conservation scientists (Longdon et al 2024). Conservation science necessarily proceeds through the accumulation of knowledge, including the efforts of multinational agencies such as the Wildlife Conservation Monitoring Centre of the UN Environmental Program, often making serious efforts to engage with local populations in the places where species are endangered, ensuring that the knowledge of those people is equitably captured and preserved, not simply privileging established Western classifications and utilisations (UNEP-WCMC 2024). Nevertheless, it is the collection and translation of this knowledge via computers that allows it to be published, prioritised, and made the focus of policy action. Despite the fact that some data collection and computation infrastructure might itself be damaging to the planet, we assume that, as in the economic logic of the UK's Resolution Foundation, productivity and increasing wealth are unquestionably good, with *more* knowledge being *better* knowledge.

I observed in my introduction above that (from my engineering standpoint) knowledge productivity involves both the accumulation and generation of knowledge. In the two years since paying considerable attention to the *accumulation* of knowledge in a contribution to SFSP III, my plan for SFSP IV is to turn instead to the question of *generating* knowledge, which might potentially be less costly than accumulating and extracting knowledge from distant forests. In particular, I want to consider the ways in which AI systems might appear to be re:generating knowledge, producing knowledge more quickly and in greater quantities, through the automated generation of text.

There are two parts to the paper, each in the tradition of *critical technical practice*, exploring AI through the use and construction of AI tools (Agre 1997). The first part reflects on the re:generation of knowledge through the writing of academic papers using AI tools, the attempts to regulate those practices, and the commercial changes in the academic publishing industry. After questioning many conventions of academic writing, the second part of the paper steps away from those conventions with autoethnographic reflection on an experimental drone metal art project, constructing a posthuman technocultural assemblage from embodied knowledge practices and artefacts. I end with an analytic comparison of these two enterprises (scientific writing and drone metal), pointing out their technical similarities while contrasting the re:generation of knowledges in each.

The Knowledge Industry of AI

Artificial Intelligence is a branch of knowledge where the *object* of study is itself the making of knowledge. This characteristic is shared, perhaps, with fields such as analytic philosophy and pure mathematics. But unlike those fields, AI has the possibly unique distinction that its practices and definitions are becoming imposed across the whole academy, wherever the practices of computer scientists are deployed as general purpose tools for knowledge work. The information processing innovations emerging from computer science research, including email, the internet, the world-wide web, laptops and word processors, have become the fundamental everyday tools underlying every field of inquiry. At a time when the latest developments of AI appear to define, or even redefine, what knowledge can possibly be, the tools produced by AI researchers are consequential for all.

This distinctive property of AI research, that it both constructs and constrains knowledge, is further complicated by the fact that recent developments in the field are led, not by universities or publicly-funded researchers, but by global corporations such as Google, Amazon, Meta and Microsoft. The advances leading to LLMs came first from corporate laboratories, drawing not only on the huge financial and computational resources already amassed by those companies, but also the resources of data that had been acquired through mechanisms of surveillance capitalism, supplemented by commercial repackaging, rebranding or outright enclosures of "open" public resources, including social media content, academic research publications, and public service initiatives such as Wikipedia. The competence of LLMs that has been achieved by encoding these many sources of public and scientific knowledge is now being promoted as if the "intelligent" LLM is itself the author of the knowledge that has been appropriated.

Academic publishers have been mobilising quickly to address threats to the business of knowledge production that are arising from LLMs. One immediate problem was the fear of an uncontrolled increase in journal submissions being automatically *authored* by AI models. The ease of producing such papers might easily overload the submission and review management processes for journal publishers. With cost of manuscript submission near zero, publishers imagined a race to the bottom, as has already happened with email spam, such that production of new academic writing would be devalued to the point of worthlessness - "enshittification" as technology commentator Cory Doctorow has named this commercial logic (2023).

A secondary potential concern is the undermining of the checks and balances within the scientific reputation economy, since these rely on the fact that the consequences of misconduct for individual academics will be sufficiently severe to act as a general disincentive to plagiarism and fabrication of results. In contrast, AI-based authors, if generally accepted as fictionally named scientist-personas, could be created or deleted like any social media troll account at the convenience of those controlling them, meaning that potential reputational harm to real individuals would no longer be the underlying sanction that governs professional science.

In order to avoid these immediate threats to their business models, most major academic publishers have come to a consensus that an AI cannot be credited as an author of a published academic paper. However, while reacting to this urgent commercial concern, publishers have also come to understand that routine use of generative models is already commonplace across the academy. A blanket ban would be unenforceable (since there is no currently effective way to test a given text as LLM-output, or prospect of any such test in the near future), and might also appear a fruitless finger-in-the-dyke, given the clear routine utility and popularity of these tools.

The response to this second challenge, again reflecting general consensus across the major publishers, has been to acknowledge that authors probably *will* be using LLMs and other text automation tools (including summarisers, paraphrasers, grammar correctors, translators, and simple sentence auto-completion). Rather than attempting to ban such widely used and useful tools, the policy of most publishers is now to acknowledge that they may be used, so long as this is declared by the author.

Re:gulating Bullshit

If regulations mandating author declarations are phrased so broadly as to include features of everyday products (e.g. Microsoft's AI *CoPilot for Office 365*), these declarations will become routine, perhaps added to every paper as has already happened with the cookie consent buttons that have been added to every website. In that case, well-meaning EU data protection regulation failed to anticipate the almost universal usage of cookies, rendering this performance of "consent" almost meaningless. The requirements of the latest European AI Act are likely to result in similarly routine warnings that text has been generated with the assistance of AI tools,

not only in academic papers, but in every newspaper edition, promotional booklet, company report, and website - perhaps alongside the cookie consent button.

I have directly asked several academic publishers to suggest a technical definition that might support their regulation of "AI". Some refer to statements from national legislative bodies or elected authorities, but these can (for example in the USA) be so broad they even apply to the heating control thermostat on my home boiler. Others refer to future prospects of AI that are apparently taken from marketing literature or science fiction scenarios, rather than any recognised algorithm. If there is no clear definitional boundary between routine functions, dangerous hazards, and science fiction, attempts at legal enforcement seem doomed to fail.

Following a televised attempt by the British Prime Minister in November 2023 to court favour from opinionated technology investor Elon Musk, I published a short blog article suggesting that, rather than regulating AI as the PM had been proposing, the UK would gain greater benefit by regulating bullshit (Blackwell 2023). I based this claim on the observation that LLM output did satisfy philosophical definitions of bullshit (cite Frankfurt), and seemed likely to bring greatest economic impact through automation of bullshit jobs (cite Graeber). On being invited to rewrite this blog post in a more scholarly register for a science policy newsletter, I gave the text to an LLM (Anthropics' Claude) with the instruction to do this for me, adding a final declaration that I had "employed multiple automated language-processing tools, including functions for spelling and grammar correction, summarisation, paraphrasing and predictive text" (Blackwell 2024).

It would be hard to say, on technical grounds, whether particular automated text tools should, or should not, be classified as "AI". My routine writing practices, including the preparation of this paper, collect text fragments typed on my phone (using predictive text), in Microsoft Word (with automated word and phrase substitutions), and online in Google Docs (which completes whole sentences). All employ text algorithms that would once have been called "AI". To prove the point, I have been writing this very paragraph on my 2020-model MacBook Pro, which suggests the next word on a dynamic Touch Bar display at the top of the keyboard. I did not need to type around half the words in the last sentence, because they were suggested to me by an algorithm, while I simply accepted each suggested word at a touch.

The automatically generated rephrasing of my bullshit blog (Blackwell 2024) did not appear in a peer reviewed journal, and the editors did not ask me for any declaration of whether I had, or had not, used AI tools. I tried to make this clear to readers, by supplementing my own declaration of automated text processing with the observation that “[t]his article is a demonstration and expansion of” the principles of bullshit I was describing. Somewhat ironically, I subsequently needed to add an acknowledgement that the “mechanically assisted text was then further adjusted with over 100 modifications made by a human copy editor,” meaning that this anonymous professional had in fact spent more time writing the piece than I had done as the credited “author”.

Re:generating Knowledge-as-text

The role of the anonymous scientific copy editor in knowledge production is particularly striking among other dramatic changes to the production of academic writing. A writer like me, inheriting the colonial privileges of writing for an international audience in my own first language of English, is able to write essays like this one fluently, whether or not assisted by AI predictive text tools. For those who may be more fluent in other languages, but are obliged to *publish* in English, copy-editing services become the gatekeepers to scientific success. Most academic publishers advertise commercial partnerships with copy-editing agencies that will turn poorly written manuscripts into publishable knowledge product⁴. Journals do not ask authors to declare the use of these services, which offer convenient sources of revenue all round, and are tacitly understood to be a necessary evil, however much they might resemble a tax on less privileged writers (Kucirkova 2023). The irony is that, while journals race to insist authors declare their own usage of AI, the copy-editing companies are also replacing human copyeditors with AI models, which are trained using the professional work of freelance contractors, but now rapidly replacing them. Editing service companies promote AI assistance software directly to graduate students who struggle with writing, but publishers are coy about which aspects of their own commercial partnerships should, or should not, be declared in the journal’s authorship policies (Blackwell and Swenson-Wright, 2024).

⁴ For example, Cambridge University Press “author services”, as with similar offerings from other academic publishers, is a Cambridge rebranding of the “Editage” product: <https://www.cambridge.org/universitypress/author-services/>

AI is thus pervading the production of academic “knowledge,” responding to policy demand for greater productivity from every nation’s scientists, and apparently resulting in a huge surge of new manuscript submissions to many established journals, especially coming from universities in larger non-English speaking countries. As academic publishers mobilise to protect their commercial remit as worldwide guarantors of scientific quality, extracting revenue from publicly-funded scientists both during the reading and writing of scientific text, they find themselves in the difficult position of having to justify the value of their own contributions, denying AI authorship, requiring authors to declare automated assistance, while simultaneously complicit in the automation of the copy-editing that was at one time the primary contribution of professional journal staff.

The remaining scientific function of an academic journal is strategic oversight of editorial decisions, and even this is being eroded. Although most journals advise their unpaid peer reviewers *not* to reveal confidential manuscripts to LLMs when writing reviews, and manuscripts submitted for professional copyediting *are* being uploaded without consent to train the publishers’ own language models, there is general silence on whether the scientific management of the journal itself should make use of AI. In a recent experiment, I responded to requests from a relatively newly established academic publisher, suggesting that they might consider using AI to reduce editorial effort. Following an invitation to join the editorial board, or to contribute as guest editor of a special issue, I responded with a special issue proposal that was generated by an LLM. The response was enthusiastic, asking me to introduce the “friend” who I said had written the text, and had suggested might be willing to serve as an editor. I carefully explained my case, that future policy for editorial work seemed to lie somewhere between scientific authorship (where AI must be declared) and copyediting (where it is not), and that the journal might therefore consider using AI in their editorial work. Unfortunately it seems that my email address has now been removed from the mailing list of this particular publisher, meaning that it has not been possible to continue the experiment.

Re:generation of Text as a Business

The open access campaigns of recent years are especially interesting, in relation to the economic function of academic publishers as the providers and producers of scientific knowledge. The open internet, world-wide web and efficient free-text search engines

created by computer scientists following the principles of open source software development presented a clear threat to the academic publishing industry, since academics were among the first to realise that their writing could quite effectively be published simply by uploading manuscripts to their own university's website. This was an existential threat to academic publishers, whose original function had been little more than distributing academic writing from individual scholars to a wider audience.

In the face of the obvious truth that nobody would pay for collation, printing and postage of content that could be more cheaply and conveniently downloaded online, journal publishers instituted a campaign claiming that public access to publicly-funded research was under threat. The greatest threat, the publishers said, resulted from the high subscription fees that they were now obliged to charge, as a result of the dwindling number of libraries still purchasing paper subscriptions. In a remarkable global consensus of poachers turned gamekeepers, scientific publishers agreed that they had always wanted their output to be available for free to anyone who needed it, and that (for example) scientists in wealthy nations and institutions ought to pay the cost of publication so that their work would be available to all. While paying to print and distribute one's own work was once derided as 'vanity publishing', the rhetoric of knowledge generation encourages research funding agencies to invest taxpayer funds on accumulating knowledge wealth, and the student debt industry subsidises citation metrics with money that could have been spent on tuition (Newfield 2018).

Academic writers supported this move to open access, apparently under the illusion that large audiences were queuing to benefit from their ideas. In reality, rather than a frustrated worldwide audience, many academics struggle to persuade their own students to read their publications. The rhetoric of open access, while supported by a logic of text-based knowledge as public good, was constructed on the basis of text as commodity, to be bundled and distributed like any mass market consumable.

The commercial and technical reconfigurations of academic publication have had perverse effects. Now that revenue is derived from authors rather than readers, profits are maximised by the number of papers written, not the number of papers read. As cost of electronic distribution approaches zero, there is no reason to restrict the size of journal issues, or the numbers of issues each year. The print issues of those journals I still receive have quadrupled in size, but many are no longer printed at all, simply deposited as thousands of unread pages in an online repository. While the rhetoric of peer review within the attention economy suggests that journal publication is a scarce

resource to be protected through rigorous double-blind quality control, in reality, through the article processing charges of open access, journals make more profit by *publishing more*, whether or not anyone reads this text.

Ever since Robert Maxwell blurred the lines between scandal sheets and science, by directing the profit-seeking instincts of the tabloid newspaperman toward academic publishing in his acquisition and massive expansion of the Pergamon (later Elsevier) empire (Cox 2002, Stevenson 2009), we have seen the conception of scientific knowledge transformed by the technological reality of Herb Simon's attention economy - that in a world with an excess of information, the scarcest commodity will be that which information consumes: human attention.

The profit opportunity spotted by Maxwell, through exploiting the scarce attention of academics, was so blatant that extracting profit from academic publication must have seemed like shooting fish in a barrel. As a newspaper proprietor, Maxwell previously had to pay his journalists for copy, pay his editors to keep the audience engaged, and make sure his product was sufficiently entertaining that people wanted to buy it. In academic publication, by contrast, the scientists write for free, the editors and peer reviewers are volunteers, and the subscriptions are bought by libraries in a captive market, where reading is prescribed as good for you, rather than purchased voluntarily for pleasure. With all elements of the publisher's business costs and revenues delivered to them on a plate, the only remaining task was to collect the profit.

I experienced this starkly, as a proud editorial board member of an Elsevier journal once so central to my PhD research that I had even bought an individual subscription. As years passed, Elsevier came to believe that this journal was under-performing, too specialised, and rejecting too many potentially revenue-generating papers. The editorial board protested that we rejected poor quality papers to protect the reputation of the journal and the scientific quality of the field. The response from Elsevier was to sack the Editor in Chief, replacing him with someone who had never published in the journal, had not attended international meetings of the field, was not known to other board members, but was willing to increase the revenue by accepting more papers for open access publication. The remaining editors were invited to resign from their own

journal, and the founder did so. Those who stayed to protest (including myself) were told that their period of service had expired and forced to leave⁵.

The business of these established publishers is now further threatened by many new entrants wanting to benefit from an opportunity so inexpensive yet profitable. The scientific community tries to sort the wheat from the chaff of open access by “impact factors” as a proxy for people actually benefiting from the knowledge being generated. However, impact factors are also constrained by the scarce commodity of human attention, so publishers are proposing “altmetrics” as evidence for readership beyond scholarly citation, with the social media statistics of TikTok or Instagram threatening to turn celebrity open-access authors into little more than tuition fee-subsidised “influencers” whose work can be liked at a click.

Cash-strapped universities are reluctant to keep paying to publish in prestigious journals, while also subscribing to buy back those very same journals for their libraries. Publishers play hard-ball when negotiating lists of journals that might qualify for discounts to either readers or authors. One interesting strategy is to shuffle titles from company to company, so that a new owner can be released from previous agreements, either reinstating publishing fees that had previously been exempt, or asking academics to instruct their universities to renegotiate terms - as happened recently to the journal *Interdisciplinary Science Reviews*, where Cambridge authors found they were explicitly excluded from open access status in the SFSP III special issue, because the journal had recently been acquired by a new publisher, leaving the authors as unwitting hostages or bargaining chips in an attempt by the company to renegotiate more favourable terms over the existing contract with Cambridge.

Re:generation as Critical Technical Practice

By juxtaposing the commercial arrangements of academic publishers with the tensions and contradictions of publishing AI-generated text, I have highlighted the ways that the production of scientific knowledge is an outcome of a global sociotechnical assemblage, in large part resulting from the economic logic of productivity growth in knowledge-as-text. That economic logic makes no fundamental distinction between the

⁵ The journal has now been renamed, but I cannot reveal the name here, since the future livelihood of my own graduate students continues to depend on publishing their research in the longest-established organ of our field.

text of science, newspapers, novels or encyclopaedias. By treating new text as an economic commodity, we have entered an era in which the automation of bullshit appears desirable. Quite transparently, the artificial intelligence of LLMs is a large-scale apparatus for the re:generation of bullshit, placing us in the quandary of how to re:ulate or even re:cognise the value of *human* intelligence as a scholarly enterprise.

An alternative to the co-option of academic writing by AI researchers can be found in the tradition of *critical technical practice* - the proposal by AI researcher turned technology commentator Philip Agre, in his “notes on trying to reform AI” (1997). Disappointed by the way that philosophical questions had become trivialised and devalued by his computer science colleagues who scarcely understood the meaning of the epistemological terminology that they claimed to be investigating, Agre suggested that engineers must become better schooled in philosophical thought. His prescription for practice also suggested that philosophers should participate in *building* AI, just as a professor of surgery, or a professor of carpentry, should hold a scalpel or use a chisel rather than simply offering opinions on what others learn through working with their hands.

I have characterised my own work as critical technical practice for years, and each of my contributions to the previous three SFSP workshops have advanced this agenda in different ways: in the first workshop, investigating the algorithmic foundations of AI in relation to the shaping of English parks and gardens, as contrasted to the forests of my home country Aotearoa New Zealand (Blackwell 2017, 2019). At the second workshop, I set out plans for investigating AI ethnographically, working with computer scientists on the African continent (Blackwell 2021, Blackwell, Damena and Tegegne 2021). The third workshop provided an opportunity to discuss the practical side of AI - the collection and processing of data - as being developed by a research student investigating conservation science with a community on the boundary of a forest reserve in Ghana (Longdon et al 2024).

At this fourth workshop, I build on an invitation I received to join a group of arts researchers exploring critical technical practice in relation to music composition and performance (Blackwell 2022). I have collaborated for years with music researchers, in the two distinct areas of live coded performance (Blackwell, Cocker et al 2022), and investigation of violin sound (Fritz, Blackwell et al 2012). Both fields draw attention to the sociocultural complex within which embodied sound worlds, craft practices and material artefacts are structured into expressive systems of signification. The second

part of this paper re:ports a very different kind of re:generation, stepping away from academic publication to investigate music and performance.

At this point in my story, it seems disingenuous to continue in a conventional academic register, especially since you (dear Reader) may be doubting whether I have written any of this myself. You might suspect me of serving up text synthesised automatically by an LLM, boosting my own academic productivity (words counted toward my “research excellence” score), while simultaneously wasting your own time and attention, in the minutes you have spent reading these words when you could more profitably have been re:generating your own.

To avoid such suspicion, I therefore turn to a different mode, of personal and embodied reflection on creative experience, extending beyond the kind of text that a disembodied LLM might re:commend for a re:search assessment exercise.

Wood, Blood, Breath and Code - embodied, posthuman and re:mixed knowledge re:generation



Figure 1. Output from StableDiffusion SDXL 1.0 prompted with "Four abstract icons to signify wood, blood, breath and code." (Overall Prompt Weight 50%, Refiner Weight 50%, Sampling method K_DPMP_2M)

Wood

Live-coding visual artist Luke Church is a photographer, programming language designer, and computational woodworker conducting computer science in the forest, based in the Highlands of Scotland, where he is re:wilding a block of commercial forest with fellow computer scientist Mariana Marasoiu (Fig. 2). But the largest block of wood in the band is Alan Blackwell's orchestral double bass, augmented with contact microphones for electronic mixing and distortion, informed both by his research into violin acoustics, and his reflections on critical technical practice in music technosystems, where he condemned the disembodiment of AI research as "Too Cool to Boogie" in a conference paper expanding on the lyrics of the 1970's disco classic *Boogie Oogie Oogie* by Taste of Honey (Blackwell 2022).



Figure 2. Forest regeneration project - Luke Church and Mariana Marasoiu

Wood meets Blood

The chainsaw, everyday equipment in Luke's forest, constantly threatens to dismantle the threshold between wood and blood. Safety is obtained with Chainsaw Trousers⁶, a phrase we appropriated to name the band. The Live Coding Manifesto proclaims: *Live coding is not about tools. Algorithms are thoughts. Chainsaws are tools. That's why algorithms are sometimes harder to notice than chainsaws.* (TOPLAP draft manifesto, February 2004). Chainsaw Trousers ironically refers to the masculine signifiers of the "trouser role" in classical opera, while admitting the challenges we face in live coding for creative education, where the girls who we'd love to inspire with new technology would rather be Billie Eilish than Billie Eilish's brother.

Blood

Vocalist Annouchka Bayley is an advocate of feminist posthumanism as a strategic orientation to her work directing a Master's programme in Arts, Creativities and Education. Her novel *The Blood Countess* is an alternative science of the past, retelling the history of Elizabeth Bathory, whose name is often appropriated in heavy metal music as little more than a sexualised vampire, but reclaimed in Bayley's novel as not only a builder of underground printing presses in defiance of the Holy Roman Emperor but quantum-entangled time traveller sharing the embodied experiences of a 21st century academic.

Blood meets Breath

Blood and breath are the bodily fluids of our cyborg ensemble. Even without the white tie and wing collar of his orchestral gigs, Alan's bald head and winter pallor suggest Count Orlock of *Nosferatu* prowling for his next meal. The band's formation celebrated Annouchka's recovery from a series of hospitalisations with an undiagnosed blood condition, resolved only in the month that *Blood Countess* emerged from the press. Meanwhile Luke was unable to sing, following an accidental ignition of chainsaw fuel, avoiding serious burns, but inhaling from a fire extinguisher. Mortal bodies haunt our narratives, drawing Chainsaw Trousers toward the body horror of David Cronenberg, or of Tsukamoto's *Tetsuo the Iron Man*, as pulse beats must be recovered within long breaths, re:generating algo-rhythms.

⁶ e.g. AT4080 - Arbortec Deep Forest Chainsaw Trousers Design A/Class 1 - Red
<https://arbortec.com/products/at4080-trouser-deep-forest-design-a-class-1-red>

Breath

The experiments of Chainsaw Trousers follow Alan's familiarity with the low bowed notes of the double bass, and Annouchka's experience learning overtone singing in Mongolia. Electronically manipulated via guitar pedals, we dramatically slow the sound of a chainsaw, pitch-shifted with the *Sonic Pi* live code language (Aaron et al 2016), and adding an archived drone performance from Thor Magnusson's *Threnoscope* (2014). These breathing systems re:mix in a foreboding soundscape that could be icelandic indie, ambient black metal, death industrial or drone rock. In this distinctive region of musical subculture (Fig. 3), we learn that these genres coalesce around the cyborg drone most familiar in Luke's Highland forest, the bagpipe.

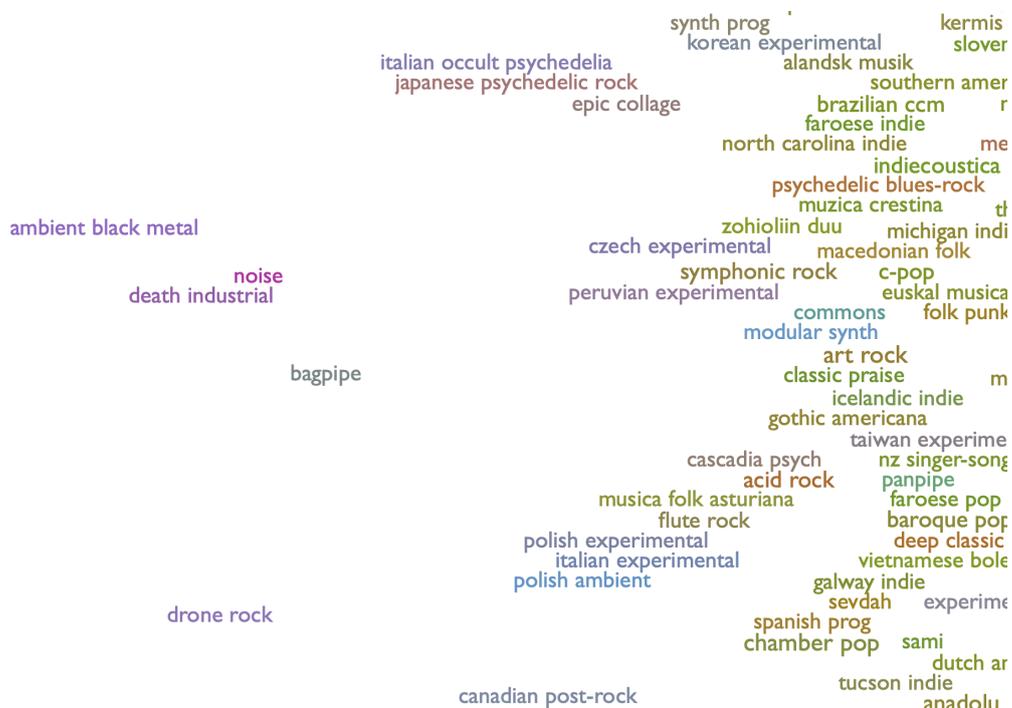


Figure 3. An extract from the interactive visualisation *Every Noise At Once* by Glenn McDonald: "an ongoing attempt at an algorithmically-generated, readability-adjusted scatter-plot of the musical genre-space, based on data tracked and analyzed for 6,291 genre-shaped distinctions by Spotify". <https://everynoise.com>, frozen 19 Nov 2023.

Breath meets Code

In hearing music, or making it, we all carve structure from sound. Whether finding the pulse within a breath, gendered pitches of vocal formants, or the harmonies in resonating wood, music is a language, stabilised, delineated and notated as one of Ingold's lines (Ingold 2007). The codes of Western music are formalised into pitch classes, key signatures, and dance times, while the rest is noise (Ross 2007). Our

experiments with drone metal explore this space of timbre, the embodied sound beyond notation that we've heard in bands like *Sunn o))), ØXN, Divide and Dissolve* or *Lingua Ignota*. When generated *from* code, these sonic textures become poetic ambiguities, demanding the computational successor metaphysics of Brian Cantwell Smith's *On the Origin of Objects* (1996), or even his father Wilfred Cantwell Smith, who denied any boundary in the transcendent: "Neither religion in general nor any one of the religions, I will contend, is in itself an intelligible entity, a valid object of inquiry or of concern either for the scholar or for the man of faith". (Smith 1962, p.12)

Code

The practice-led research of Chainsaw Trousers finds knowledge in the body, extracting abstraction from a sequence of sound. Our explorations of wood, blood, breath and code are re:mixes of craft assemblage, and the performance of personal identity and experience. In practice, we step aside from the preoccupations of knowledge as text, with a posthumanist analytic frame attending to other ways of knowing. As live coders, we take abstract notation as a resource for improvisation, an immaterial material. In the era of generative machine learning systems, we re:claim pastiche as a tool, understanding alternative knowledge production in the light of re:mix culture. These practices of textu(r)al collage re:lect the mediated distributed agency of many creators, engineers, coders, luthiers, and songwriters whose own knowledge has become embodied in the tools and materials we work with.

Code meets Wood

Michael Castelle (2020) has observed that deep neural networks operate like an array of electric guitar pedals, alternating layers of distortion and filtering. Just as Chainsaw Trousers constructs new sounds from chains of components re:mixing our own wood and breath, so the neural network generates textures from the matrix or humus of conceptual origination, derived from multiple modes of embodied knowing, extending far beyond knowledge as text. Drawing is in the eye, not the pencil (Edwards 1979), and music happens in the ear, not in the violin (Cross 2009). The internet transformed reading before it transformed writing, and we understand that creative research is in the question, not the answer.

Stripes of Causation / Tools for our Times

This paper has presented reflections on science in the forest and science in the past, attending to the tools that scientists use in their work, and specifically to knowledge tools - the electronic/algorithmic processing of text and data popularly described as “machine learning” and “artificial intelligence”. We are still coming to understand how the mechanical calculation of statistical regularities relates to human experiences of learning, thinking and knowing that are considered central to the practice of science. But in a sense, these statistical machines *are* fundamentally relevant to pasts and to forests. They are immense extrapolators, trained with data collected across the internet, from traces of our own actions, including archives of discourse in all registers and degrees of seriousness and intent. From that stabilised archive of the past, they presume to act re:peatedly in the future. Machine learning systems are also immense simplifiers, condensing the blooming buzzing forest of the world’s lives into a statistical model that regularises and abstracts the patterns to be re:generated.

The scale and success of these strategies has brought AI to the attention of the whole world, to an extent that some suggest science itself will be transformed by a new paradigm of production at scale (Bishop 2022). Information theorists observe that the human brain is an information processing system, encoding patterns from observation, no more able than any physical information system to output anything else than a functional transformation of what has been input. Human science has always generated knowledge from the past of scholarship and from the forest of observational experience, so might the mechanical re:generation of knowledge represent a new crisis of free will, as these tools make it seem that not only the physical actions of our bodies but the products of our minds can be mechanically explained and simulated?

In the current work of my research group (when not performing drone metal), we develop Bayesian representations of causality as a basic tool for a new kind of literacy, where theories of knowledge and discourse can be related to mechanical data processing, serving as the starting point for more intentional design of intellectual automation. The educational experiments we are undertaking determinedly engage with knowledge systems beyond the European and Western tradition, including indigenous knowledge of sub-Saharan Africa, and of the Pacific Moana Oceania (Lythberg et al forthcoming). This is a challenging project, resisting the hegemonic cultural infrastructure and legacies of the Anglo-American attention economy, as

illustrated in the “unconsciously” biased hallucinations of generative AI models (e.g. Fig 4).



*Figure 4. Output from NightCafe CKTP model prompted with "Time is a tiger inside a forest of lives, wielding stripes of causation on its back." from Annouchka Bayley, *The Blood Countess*. Leicester: Troubador (2024), p.147.*

My own proposal is that more literate and critical engagement with these tools will re:orient populations and institutions of learning toward technical opportunities, while enabling and prioritising the embodied cultures and creative experiences that humans most value. Automated text technologies can offer much more than faster generation of more knowledge, if they support re:generation - an information-theoretically nutrient pastiche that nourishes growth, even where that fertiliser might be difficult to distinguish from bullshit. Pastiche and bullshit are everyday realities of the human condition, but the purpose of science is surely to distinguish them.

Rather than neo-Luddite rejection of new automation tools, a more likely strategy for productive transformation is to embrace and understand the value of pastiche in the light of creative remixing, already central to digital arts and music, where the user of information technologies is valued above all as a listener, a conscious agent attending to many sources, both immersed in and constructing the technocultural craft assemblage of embodied human knowledge.

References

Aaron, S., Blackwell, A.F. and Burnard, P. (2016). The development of Sonic Pi and its use in educational partnerships: co-creating pedagogies for learning computer programming. *Journal of Music, Technology and Education* 9(1), 75-94.

Agre, Philip. E. "Towards a Critical Technical Practice: Lessons Learned in Trying to Reform AI". In *Social Science, Technical Systems and Cooperative Work: Beyond the Great Divide*, edited by Geoffrey Bowker, Susan Leigh Star, Les Gasser and William Turner, 131-157. Mahwah, NJ: Lawrence Erlbaum, 1997.

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. "On the dangers of stochastic parrots: Can language models be too big? 🦜." In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pp. 610-623. 2021.

Bishop, Christopher. AI4Science to empower the fifth paradigm of scientific discovery (July 7, 2022). <https://www.microsoft.com/en-us/research/blog/ai4science-to-empower-the-fifth-paradigm-of-scientific-discovery/>

Blackwell, A.F. (2017). Objective Functions, Deep Learning and Random Forests. Contribution to the workshop Science in the Forest, Science in the Past, convened at the Needham Research Institute, Cambridge.
<https://www.cl.cam.ac.uk/~afb21/publications/Blackwell-ObjectiveFunctions.pdf>

Blackwell, A.F. (2019). Objective Functions: (In)humanity and Inequity in Artificial Intelligence. *HAU: Journal of Ethnographic Theory* 9(1), 137-146.

Blackwell, A.F. (2021). Ethnographic artificial intelligence. *Interdisciplinary Science Reviews* 46(1-2), 198-211, DOI: 10.1080/03080188.2020.1840226

Blackwell, A.F. (2022). Too Cool to Boogie: Craft, culture and critique in computing. In J. Impett (Ed.) *Sound Work: Composition as Critical Technical Practice*. Leuven University Press / Orpheus Institute, pp. 15-33.

Blackwell, A.F. (2023). Oops! We automated bullshit. Blog post, 9 November 2023.

Blackwell, A.F. (2024). *Moral Codes: Designing alternatives to AI*. MIT Press.

Blackwell, A.F. (2024). What are Large Language Models Good For? NORRAG Policy Insights. NORRAG: Network for International policies and cooperation in education.

Blackwell, A.F., Damena, A. and Tegegne, T. (2021). Inventing artificial intelligence in Ethiopia. *Interdisciplinary Science Reviews*, 46(3), 363-385. DOI: 10.1080/03080188.2020.1830234

Blackwell, A.F., Cocker, E., Cox, G., Magnusson, T. and McLean, A. (2022). *Live Coding: A User's Manual*. Cambridge, Mass.: MIT Press. Publisher site (open access): <https://doi.org/10.7551/mitpress/13770.001.0001>

Blackwell, Alan and Swenson-Wright, Zoe (2024). Editing companies are stealing unpublished research to train their AI. *Times Higher Education*, January 12, 2024.

Castelle, Michael. Methods training seminar for the Mellon-Sawyer seminar Histories of AI, October 2020 (online).

Cox, Brian. "The Pergamon phenomenon 1951–1991: Robert Maxwell and scientific publishing." *Learned Publishing* 15.4 (2002): 273-278.

Crawford, Kate. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press, 2021.

Cross, Ian. (2009). The evolutionary nature of musical meaning. *Musicae Scientiae*, 13(2_suppl), 179-200. <https://doi.org/10.1177/1029864909013002091>

Doctorow, Cory. The 'Enshittification' of TikTok or how, exactly, platforms die. *Wired Magazine*, January 2023. <https://www.wired.com/story/tiktok-platforms-cory-doctorow/>

Edwards, Betty (1979). *Drawing on the right side of the brain: a course in enhancing creativity and artistic confidence*. Los Angeles; New York: J. P. Tarcher; distributed by St. Martin's Press.

Fritz, C., Blackwell, A.F., Cross, I., Woodhouse, J. and Moore, B. (2012). Exploring violin sound quality: Investigating English timbre descriptors and correlating

resynthesized acoustical modifications with perceptual properties. *Journal of the Acoustical Society of America* 131(1):783-94.

Ingold, Tim (2007). *Lines: A Brief History*. Routledge, Oxon, UK.

Kucirkova, Natalia. AI writing tools will not fix academia's language discrimination problem. *Times Higher Education Supplement*, September 5, 2023

Longdon, J., Gabrys, J. and Blackwell, A.F. (2024). Taking Data Science into the Forest. *Interdisciplinary Science Reviews*.

Lythberg, Billie, Rachel Maunganui Wolfgramm, Albert L. Refiti, Alan F. Blackwell
Making relations: reimagining AI through crafted and embodied knowledge systems in Moana Oceania. Manuscript in preparation for *Big Data & Society*.

Magnusson, Thor. "Improvising with the Threnoscope: Integrating Code, Hardware, GUI, Network, and Graphic Scores." In *Proc. New Interfaces for Musical Expression (NIME)*, pp. 19-22. 2014.

Newfield, Christopher. *The great mistake: How we wrecked public universities and how we can fix them*. Johns Hopkins University Press, 2018.

Stevenson, Iain. "Robert Maxwell and the invention of modern scientific publishing." *Publishing History* 65 (2009): 97-113.

Smith, Brian Cantwell *On the Origin of Objects*. MIT Press (1996)

Smith, Wilfrid Cantwell. *The Meaning and End of Religion*, New York: Macmillan (1962), p.12

Rishi Sunak & Elon Musk: Talk AI, Tech & the Future. Video interview.
<https://youtu.be/R2meHtrO1n8?si=iXWzThqa7b4YBNys>

Ross, Alex. 2007. *The Rest is Noise: Listening to the Twentieth Century*. New York: Picador.

UN Environment Programme World Conservation Monitoring Centre (UNEP-WCMC).
Leading Database on Conservation by Indigenous Peoples and Local Communities
Opens to Public. Blog post, Jan 2024.

<https://www.unep-wcmc.org/en/news/leading-database-on-conservation-by-indigenous-peoples-and-local-communities-opens-to-public>

Vonnegut, Kurt. *Galápagos: A Novel*. New York: Delacorte Press, 1985.