

Intention Is All You Need

Advait Sarkar

Microsoft Research, University of Cambridge, University College London
advait@microsoft.com

Abstract

Among the many narratives of the transformative power of Generative AI is one that sees in the world a latent nation of programmers who need to wield nothing but intentions and natural language to render their ideas in software. In this paper, this outlook is problematised in two ways. First, it is observed that generative AI is not a neutral vehicle of intention. Multiple recent studies paint a picture of the “mechanised convergence” phenomenon, namely, that generative AI has a homogenising effect on intention. Second, it is observed that the formation of intention itself is immensely challenging. Constraints, materiality, and resistance can offer paths to design metaphors for intentional tools. Finally, existentialist approaches to intention are discussed and possible implications for programming are proposed in the form of a speculative, illustrative set of intentional programming practices.

1. The “Intention Is All You Need” Picture of Programming with Generative AI

What is programming? Blackwell’s succinct and influential definition is that programming is any activity exhibiting the property “*that the user is not directly manipulating observable things, but specifying behaviour to occur at some future time*” (Blackwell, 2002). Behaviour is specified through an interface, commonly a notation, which we call a programming language. Therein lies the source and objective of all research in the psychology and design of programming: the study of the use and improvement of the interfaces, notations, and languages for specifying behaviour.

The value of such study is called into question with the introduction of Generative Artificial Intelligence (GenAI), which can be defined as any “*end-user tool [...] whose technical implementation includes a generative model based on deep learning*”.¹ GenAI captures the relationships between natural language specifications of behaviour, and the translations of that behaviour into programming notation, implicit in enormous training datasets. The power of translation thus captured can be stochastically replayed on demand (Blackwell, 2020). What could this mean for research in the user-centred design of programming languages? One perspective anticipates nothing less than its obsolescence:

“The programming barrier [with GenAI] is incredibly low. We have closed the digital divide. Everyone is a programmer now - you just have to say something to the computer.”²

“Up until now, in order to create software, you had to be a professional software developer. You had to understand, speak and interpret the highly complex, sometimes nonsensical language of a machine that we call code. [...] But with GenAI] We have struck a new fusion between the language of a human and a machine. With Copilot, any person can now build software in any human language with a single written prompt. [...] going forward, every person, no matter what language they speak, will also have the power to speak machine. Any human language is now the only skill that you need to start computer programming.”³

¹There is no definitional consensus on this term. A rationale for the definition adopted here is given by Sarkar (2023d).

²Huang, 2023. (Source: <https://www.reuters.com/technology/ai-means-everyone-can-now-be-programmer-nvidia-chief-says-2023-05-29/>, accessed July 2024.)

³Dohmke, 2024. (Source: https://www.ted.com/talks/thomas_dohmke_with_ai_anyone_can_be_a_coder_now, accessed July 2024.)

“Since the launch of GPT-4 in 2023, the generation of whole apps from simple natural language requirements has become an active research area. [...] Our vision is that by 2030 end users will build and deploy whole apps just from natural requirements.”⁴

“Programming will be obsolete. [...] the conventional idea of ‘writing a program’ is headed for extinction [...] all programs in the future will ultimately be written by AIs, with humans relegated to, at best, a supervisory role. [...] The engineers of the future will, in a few keystrokes, fire up an instance of a four-quintillion-parameter model that already encodes the full extent of human knowledge (and then some), ready to be given any task required of the machine.”⁵

The promise of GenAI for programming, therefore, is to transform programming into an activity where expertise in specialised notations and languages for specifying behaviour are unnecessary. One merely has to say what one wishes the program to do, and GenAI does the rest. The interaction design challenges of programming are solved.⁶ Intention is all you need.

There are many problems with this picture. There are compelling reasons for continuing to engage with formal notations, even and perhaps especially when GenAI is in play (Sarkar, 2023d). Moreover, language in general, and the language of prompts used to direct GenAI in particular, is most certainly not a flawless, transparent route for the expression of intent. Johnny can’t prompt (Zamfirescu-Pereira et al., 2023). Johnny can’t figure out what level of abstraction to write his prompts in, either (Liu et al., 2023; Sarkar et al., 2022). Thinking about prompting is hard for Johnny, and thinking about thinking about prompting is hard, too (Tankelevitch et al., 2024). Prompting “dialects” might evolve in much the same way as specialised uses of natural language do in domains such as scientific and legal communication, through disciplinary norms and professional consensus, and to acquire such language will require users to undergo analogous processes of disciplinary and professional acculturation (Sarkar, 2023d). But these problems are not the primary concern in this paper.

There is a rather more fundamental pair of problems with the idea that intention is all you need (to program with GenAI): it assumes that GenAI does not interfere with intention. Moreover, it takes for granted that intentions are easy to form. Both premises will be questioned in turn.

2. Mechanised Convergence: The Homogenising Effect of AI on Intention

Contrary to not interfering with intention, AI supplies intention. It does so in a way that can be described as *mechanised convergence* (Sarkar, 2023b), drawing on Walter Benjamin’s concept of mechanical reproduction (W. Benjamin, 1935). Mechanised convergence describes the idea that the automation or mechanisation of work leads to a convergence in the space of outputs. Standardisation is necessary for factory logic to function. For a machine to be repeatable at speed, its inputs and outputs need to be repeatable at speed, too. You can have any colour as long as it’s black.

Here is some of the evidence that GenAI has a mechanised convergence effect:

- Predictive text encourages predictable writing (Arnold et al., 2020). In an image captioning task, when participants use predictive text entry systems, captions written with suggestions are shorter and use fewer words than the system does not predict. A similar effect occurs in identifier names when programmers use a GenAI tool such as GitHub Copilot to assist them in writing code (Lee

⁴Robinson et al. (2024)

⁵Welsh (2022)

⁶One is reminded of similar claims made during the early days of spreadsheets or about any number of visual programming languages. E.g., Benjamin Rosen, a PC industry analyst for Morgan Stanley, later a key funder of Lotus and Compaq, noted in 1979 that *“In minutes, people who have never used a computer are writing and using programs [...] the user need not know anything about computers or programming in order to derive Visicalc’s benefits. You construct a Visicalc program much as you would define a problem on a sheet of paper or a blackboard”* (Rosen, 1979).

[et al., 2024](#)). This effect occurs even when the suggestions are merely visible and not actionable (i.e., cannot be accepted using a keyboard shortcut).

- Similarly, a large study (n=293) of participants writing short stories with varying degrees of AI assistance found that exposure to GenAI “ideas” leads to a reduced diversity of content ([Doshi & Hauser, 2023](#)). Participants exposed to even a single GenAI suggestion produce stories similar to the average of the other stories in the same experimental condition.
- A large study (n=758) of strategy consultants at BCG examined the effects of ChatGPT use on a set of consultancy tasks ([Dell’Acqua et al., 2023](#)). The majority of participants with access to ChatGPT retain a very high amount of its response – typically around 90% – in their submitted work. Participants without access to ChatGPT produce ideas with more conceptual variation than those with access, showing that usage of ChatGPT reduces the range of ideas generated. The variation across responses produced by ChatGPT is smaller than what human participants produce on their own.
- Large language models have a “homogenization effect” on creative ideation ([Anderson et al., 2024](#)). In a creative ideation task, participants produce less semantically distinct ideas when using ChatGPT. Moreover, participants feel less responsible for ideas produced with ChatGPT assistance.
- A large study (n=115) finds that conversational search built on GenAI increases selective exposure compared to conventional search ([Sharma et al., 2024](#)). Users engage in more biased information querying with conversational search, and the bias is exacerbated when the model is itself “opinionated” to reinforce the user’s views. The authors call this a “generative echo chamber”.
- Similarly, a large study (n=1506) of co-writing with GenAI found that using an “opinionated” language model affects the opinions expressed in participants’ writing and moreover, actually shifts their opinions as measured in a subsequent attitude survey ([Jakesch et al., 2023](#)). A related effect, termed “drifting”, has been observed in novice programmers, where the tendency to accept and adapt code generated by the system leads programmers away from a correct solution ([Prather et al., 2023](#)).

Mechanised convergence signals an odd reversal (or perhaps intensification) of Dennett’s “intentional stance” ([Dennett, 1971](#)), wherein we not only ascribe intention to these systems but also delegate it, sometimes wilfully, other times unknowingly.

The intention supplied by GenAI through mechanised convergence has a complex source, combining influences of its training data, and the biases and heuristics encoded by the system developers. However at its core, mechanised convergence is the ultimate outcome of the old statistical logics of uncovering underlying natural “laws” ([Blackwell, 2020](#); [Sarkar, 2023a](#)). The statistical machine eliminates “noise” (diversity) to predict “signal” (uniformity). The statistical machine is the triumph of the Enlightenment aesthetic faith in nature’s having an underlying elegance or simplicity that is obscured from view by imperfect forms. It should come as no surprise that machines that are built to search for Platonic ideals reflect back to us a mechanically converged picture of the world, making quiddity of haecceity.

It is important to note that the effect on intent as demonstrated in these studies is an *aggregate tendency* that likely does not square with individual phenomenal perceptions of GenAI use. At the granularity of individual interactions, the experience of GenAI might well be as a passive translator, not active supplier, of intent. The nudge towards standardised, centralised, averaged, generic, and statistically optimised answers may be barely perceptible. Yet the data demonstrates that these nudges in fact have a measurable cumulative effect on knowledge work.

As Winner sets out, artefacts have politics ([Winner, 1980](#)). The design features of a technology enable certain forms of power, and the decision to adopt a particular technology requires certain power relations to be enacted. Putting it in Winner’s terms, convergence is the politics of AI, the artefact.

As McLuhan sets out, the medium is the message (McLuhan, 1964). There is an effect of a particular medium, be it typography, radio, or television, on the human sensorium that is quite distinct from any particular content being conveyed through that medium. The effect of the medium overwhelms the content and makes it incidental. Putting it in McLuhan's terms, convergence is the message of AI, the medium.

McLuhan predicted that electric technology and programmability would reverse the convergence tendencies of factory logic. He gives the example of a programmable tailpipe machine: *“A new automatic machine for making automobile tailpipes [...] starting with lengths of ordinary pipe, it is possible to make eighty different kinds of tailpipe in succession, as rapidly, as easily, and as cheaply as it is to make eighty of the same kind. And the characteristic of electric automation is all in this direction of return to the general-purpose handicraft flexibility that our own hands possess. The programming can now include endless changes of program.”*

Taken to its logical conclusion, McLuhan makes a claim that is strikingly similar to the narrative that intention is all you need: *“the older mechanistic idea of “jobs,” or fragmented tasks and specialist slots for “workers,” becomes meaningless under automation. [...] The very toil of man now becomes a kind of enlightenment. As unfallen Adam in the Garden of Eden was appointed the task of the contemplation and naming of creatures, so with automation. We have now only to name and program a process or a product in order for it to be accomplished. Is it not rather like the case of Al Capp's Schmoos? One had only to look at a Schmoos and think longingly of pork chops or caviar, and the Schmoos ecstatically transformed itself into the object of desire. Automation brings us into the world of the Schmoos. The custom-built supplants the mass-produced.”* As we have seen, the vast programmability of GenAI does not necessarily result in a *“return to [...] general-purpose handicraft flexibility”*, rather, it has enabled a newer, subtler, and more pervasive form of the *“fragmentalized and repetitive routines of the mechanical era”*. Through the mechanised convergence of knowledge work through GenAI, the principle of interface design becomes WYGIWYG – What You Get Is What You Get.

Postman, who builds on McLuhan, more accurately reappraised the effect of the electric age on intention (Postman, 1985). He explains that the information age has resulted not in an Orwellian dystopia where intentions are surveilled and constrained, but rather a Huxleyan one, where intentions are numbed: *“What Orwell feared were those who would ban books. What Huxley feared was that there would be no reason to ban a book, for there would be no one who wanted to read one. Orwell feared those who would deprive us of information. Huxley feared those who would give us so much that we would be reduced to passivity and egoism. Orwell feared that the truth would be concealed from us. Huxley feared the truth would be drowned in a sea of irrelevance. Orwell feared we would become a captive culture. Huxley feared we would become a trivial culture [...]”*. We inhabit not Foucault's society of discipline (Foucault, 1977; O'Neill, 1986), but Deleuze's society of control (Deleuze, 1992).

This scenario is undesirable, not least because mechanised convergence implies a reduction in the rate at which new ideas are generated, and an increase in repetition and replay of existing ideas. What kind of culture springs from the consumption and emission of an increasingly convergent set of increasingly recycled ideas? A derivative, “stuck” culture, is the diagnosis of technology critic Paul Skallas.⁷ Even for GenAI itself, the indications are that the roads of autophagy lead to madness; the roads of recursion lead to cursed collapse (Alemohammad et al., 2023; Shumailov et al., 2024; Bohacek & Farid, 2023; Gerstgrasser et al., 2024).

Mechanised convergence, as a tendency of automation more broadly, creates a crisis of intentionality: a culture that has lost the capacity to intend, does not realise, and does not care.

⁷<https://lindynewsletter.beehiiv.com/p/culture-stuck>, accessed July 2024. Related is the concept of “refinement culture”; *“Refinement culture can be summarized as a general streamlining and removal of any unique characteristics. Refinement Culture is one dimensional and removes variety from the environment. It's optimized.”* <https://lindynewsletter.beehiiv.com/p/refinement-culture>, accessed July 2024, <https://medium.com/@lindynewsletter/refinement-culture-51d96726c642>, accessed 2024.

3. Interlude: Babbage’s Intentional Programmer

Describing what GenAI does to intention as a “crisis” implies that we need to do something about it. Indeed, what we need to do about it is to promote the active cultivation of the capacity to intend.⁸

Since this is PPIG, we can start by considering the intentions of programmers. What the tendency for mechanised convergence tells us is that, prior to specifying behaviour, programming must be about forming an intention for behaviour. A definition of programming that centres intention, rather than specification, evokes a rather older philosophy of programming that we can draw from the crisis in theology at the time of Babbage.

Science (more precisely, natural philosophy) in post-Enlightenment Britain at the time of Babbage was grappling with the apparent contradiction of divine miracles – acts of God outside the laws of nature created by God – which Hume had famously argued could not be rationally supported (Hume, 1748). In aiming to discover mathematical laws such as those of Newton, which could accurately describe and predict nature, natural philosophers operating within the frameworks of Deism and Christianity struggled to reconcile their work and faith.

Babbage found in his Difference Engine the possibility for reinterpreting miracles as part of the natural divine order. Using a “feedback mechanism” that connected two gear wheels, Babbage was able to encode programs that, after a certain number of iterations, would change their behaviour. For example, he would demonstrate a program that counts the integers 1, 2, 3 ... up to 100, at which point the program would change and start counting in steps of two: 102, 104, 106 ... etc. In demonstration-sermons delivered to rapturous audiences, he used this example to explain his theory of God as a *divine programmer* (Snyder, 2012). A miracle was thus explained as a shift in the program. God’s intervention to perform apparent miracles was not an aberration against universal, constant laws – it was merely the manifestation of a deeper and misunderstood universal law, a deeper plan, a deeper intention.

It is instructive that Babbage’s conception of programming and intention centred around shifts, or deviations, from the expected. A machine that continues to execute the same predictable behaviour is not a program, it is simply a machine. It is in the departure from convergent behaviour that evidence of programming emerges as activity and divinity. For Babbage, to converge is human, to deviate divine. To execute is human, to program divine. To specify is human, to intend divine.

4. Sources of Intention: Constraints, Materiality, and Resistance

Returning to our objective – to promote the active cultivation of the capacity to intend – it is worth briefly exploring a few perspectives on the sources of intentions.

Much intention appears to arise as a result of interaction with the external world. Practitioners of creative arts and research in creativity have long noted the role of constraints in shaping and facilitating creativity (Stokes, 2005; May, 1975). Materiality and resistance are essential to craftsmanship; any material, by virtue of its properties and resistances, participates in an ongoing dialogue with the craftsman’s intentions (Basman, 2016). According to material engagement theory⁹ (Malafouris, 2019), “*Our forms of bodily extension and material engagement are not simply external markers of a distinctive human mental architecture. Rather, they actively and meaningfully participate in the process we call mind*”. As such, the role of material as a source of intention can be seen as a form of extended cognition, or at the very least external cognition (Turner, 2016), notwithstanding challenges to these ideas (Rupert, 2004).

A sculptor must consider how pliable or fragile their material is, what tolerances and fine details can be accomplished, how gravity will constrain the scale and orientation of their figures. A carpenter must consider the grain of their wood, where cuts and incisions can be made. A painter using watercolours must consider and exploit the additive translucency of that medium, one using oils must consider the opacity of theirs. It is telling that the archetypical dimension in the Cognitive Dimensions of Notations

⁸Much as R. Benjamin (2024) calls for us to cultivate the capacity to imagine.

⁹Thanks to Ava Scott for identifying this connection.

(Green, 1989) is *viscosity*, a metaphor rooted in materials and resistances, aiming to bridge them with the seemingly immaterial and disembodied world of notations.

Some intentions even rejoice in the contradiction of others: for example, the objective of subversive gameplay styles is to ignore the received goals of the game and invent one's own (Flanagan, 2009), it is playing the infinite game whose objective is to continue playing, not the finite game whose objective is to win (Carse, 1986). Solving the continuous puzzles posed by these resistances, having a vision pushed, pulled, and evolved, is the pleasure and intentionality of craftsmanship. These are not destructive resistances that hinder the realisation of an intention; they are productive ones that facilitate it.

Exploratory programming (Kery & Myers, 2017) exemplifies how the materialities and resistances of programming are exploited to shape intention. In exploratory programming, the programmer's goal is unknown or ill-defined. The objective of the process is to discover or create an intention, to formulate a problem. The formulation of a problem co-exists with and cannot be separated from its solution (Rittel & Webber, 1973; Sarkar, 2023c). This is also the case in the end-user programming activity of interactive machine learning, or interactive analytical modelling (Sarkar, 2016b), where the goal is ill-defined and the objective is to create one, through a constructivist loop of interaction between ideas and experiences (Sarkar, 2016a).

There have been proposals to design GenAI systems that introduce productive resistances as catalysts for the development of intention. Rather than an assistant, AI can act as a critic or provocateur (Sarkar, 2024; Sarkar et al., 2024). AI can be antagonistic (Cai et al., 2024). AI can cause cognitive glitches (Hollanek, 2019). AI can act as cognitive forcing functions (Buçinca et al., 2021). These proposals are counter to traditional narratives of system support, system disappearance, and system non-interference. They can be seen as successors to previous counternarratives raised by researchers such as critiques of the doctrines of simplicity and gradualism (Sarkar, 2023c), critiques of seamlessness (Chalmers & MacColl, 2003), critiques of reversible interactions (Rossmly et al., 2023), the case for design frictions and microboundaries (Cox et al., 2016), reframing of ambiguity as design resource (Gaver et al., 2003), and calls for attention checks in AI use (Gould et al., 2024).¹⁰

The concept of resistance could be key to framing the design objectives for intentional GenAI tools. Our current explorations of improving critical thinking with GenAI (e.g., Sarkar et al. (2024)) are strictly *additive*: let's augment AI interaction and output with prompts, text, visualisations, etc. that get the user thinking. However, this approach increases the cognitive burden by asking users to consume and reflect on more information. We know that people don't always enjoy, or want, more information. Particularly when it comes to the user interfaces of discretionary software, they usually want less (Carroll & Rosson, 1987; Sarkar, 2023c). The additive approach may be starting by fighting a losing battle, one in which we try to design the smallest, most stimulating, most rewarding "consumable" that creates user reflection, without incurring undesirable attentional costs. The idea of resistance provides a different starting point. How can we build GenAI tools with inherent, productive resistances that are part of working with the tool, not an additional thing that users need to "pay" attention to? How can the experience of resistances in the interface feel more like the pliability of clay, or the translucency of paint? This is an open avenue for future work.

¹⁰It is worth observing that while such counternarratives often involve calls for greater, more critical, and more reflective user engagement and participation with technology, it should not be assumed that intentionality always entails participation or action. Observation itself is not intrinsically passive. This point is well made by Pfaller (2017): "*Two philosophical premises silently played a decisive role in this triumphal march of participation: first, the idea that the relation between transmitter and receiver represents a hierarchy and that the elimination of this hierarchy therefore has to amount to a democratisation; and secondly, the idea that it is more desirable for spectators to participate than to spectate [... however;] the relation between transmitter and receiver does not always represent a hierarchy. And when it does, then it is not always in favour of the transmitter [...] This is why it is misleading to believe that activating the audience in art is automatically and always tantamount to their liberation. Because could not the exact opposite be the case: could the enthusiasm for joining in produced by art not deprive people of the necessary refractoriness that they would need in political life in order not to be immediately enthused by every neoliberal or reactionary or even fascist appeal to 'actively' join in, and pursue this with a feeling of liberation?*"

5. Existentialist Approaches to Intention

So far we have been considering intention at relatively small scale: instances of knowledge work and GenAI use. But intentions, like goals, form hierarchies. Intentions are not isolated and independent, they are related and convergent. To what do they converge? At this point we shall make a somewhat abrupt leap outwards and consider the most expansive scope of intention – as enacted over the course of an entire life.

An evolutionary account might attempt to trace human intentions back to fundamental physiological concerns: we form intentions to continue survival, to avoid fear, to ensure comfort, to maximise pleasure, to minimise pain. These can certainly account for some intentions. The concept of intention has much in common with free will – loosely defined, one’s capacity to act differently to how one did, in fact, act. Free will is not the same as intention, but it can be viewed as a precondition for true intention. Neuroscientific work purporting to demonstrate (a lack of) free will has been criticised by philosophers because (among other objections), we do not have a suitably good picture that connects short-term choices dominated by low-level psychological phenomena (such as choosing to push the left button or the right button) to the complex, long-term, highly planned and goal-oriented intentions (such as the intention to commit a crime) that pose the truly consequential ethical challenges to free will (Mele, 2019). The evolutionary account is part of a broader category of *teleosemantic* theories of intention (Jacob, 2023) according to which design (evolutionary or artificial) supplies a function (τέλος), which in turn supplies intention.

In considering whether human intention can truly be reduced to evolutionary or functional needs, I am drawn to the argument made by feminist anthropologist Payal Arora in her closing keynote for the 2022 CHI conference (Arora, 2022). She criticizes Maslow’s famous hierarchy of needs that places physiological and safety needs at the bottom, rising to esteem and self-actualisation at the top. The conventional reading is that needs at the bottom of the pyramid need to be satisfied, the foundation of the pyramid needs to be built, before one can proceed to the higher levels. This is a fairly influential way of thinking and often dictates the way in which social aid and rescue efforts are prioritised: focus on food, water, and shelter first, and joy, play, growth, education, and dignity later. Arora finds that this picture does not correspond with her observations in her extensive ethnographic work with precarious, oppressed, and underprivileged groups. Instead, she proposes that the pyramid is upside down. What she finds is that self-actualisation is what people need first, and are willing to sacrifice safety needs to get it. People leave secure work when the nature of that work threatens their dignity, even if this places them in financial hardship. People leave homes where they cannot express their identity, or are not accepted for who they are, even if this might leave them without a roof over their head. A line from the poet James Oppenheim captures the sentiment:

*“Our days shall not be sweated from birth until life closes —
Hearts starve as well as bodies: Give us Bread, but give us Roses.”*

If not entirely upside down, then at the very least Maslow’s hierarchy is not a unidirectional ladder to climb, but a set of considerations and influences that are continually negotiated and traded-off. Physiology and evolution are part of intention formation, but far from the entire picture. Where can we look for a perspective on intention that aligns with Arora’s observations? Moreover, is there an approach that not only identifies the source of intention, but prescribes a method for cultivating it?

Elaborating the consequences of the idea that the active cultivation of intention is *the* core virtue in an inherently meaningless world is precisely the project of existentialist philosophy.

The absence of any inherent purpose to life is the starting point. Per Sartre (1943), “existence precedes essence”; individuals first exist without purpose and must subsequently forge their essence, or identity, through their actions. Angst, or existential anxiety, arises from the realization of one’s freedom and the infinite possibilities it entails (Kierkegaard, 1844). Existentialists see angst as a motivator rather than an obstacle.

Authenticity is one expression of existentialist intention. It is the pursuit of living in accordance with one's true self and values, rather than conforming to societal norms, and is essential for genuine existence (Heidegger, 1927). Authenticity requires a conscious effort to understand and act upon personal convictions, even in the face of adversity or societal pressure (Kierkegaard, 1843; de Beauvoir, 1948). Other sources of intentionality, besides authenticity, go beyond the individual. Kierkegaard's (Kierkegaard, 1849) "leap of faith" suggests that to escape from existential despair requires acknowledging the limits of rational reflection and an individual's relationship with the divine. Moreover, to seek engagement with the world is to step beyond oneself, to interact with others, and to find and create meaning through these actions (Jaspers & Saner, 1932). Similarly, de Beauvoir (1948) points out that our individual subject-like freedom is complemented by an object-like unfreedom ("facticity"), deriving an ethics of freedom that advocates for actions that respect the freedom of others.

Camus (1942) counsels individuals to accept "the absurd": the tension between the human search for meaning and a universe that is silent in response, to recognize the lack of inherent meaning in the world and to take on the task of creating their own purpose. Camus rejects "solutions" to the absurd proposed by prior philosophers, such as Kierkegaard, as "philosophical suicide". To Camus, seeking overarching meaning despite the absurd is seeking to resolve, minimise, sidestep, or ignore the absurd, not acknowledging it.

Camus rejects a forced imposition of meaning where there is none. A leap of faith is a form of escape. Incidentally, a forced imposition of meaning is precisely the *modus operandi* of GenAI: for language to be produced by arithmetic means it is necessary to encode language in a uniform, rational vector space. Sense and nonsense alike are thus enumerated and made commensurable. *King–Man+Woman=Queen* (Mikolov et al., 2013). Before carefully designed guardrails (themselves a form of escape) made it more difficult to do so, it was easy to elicit answers to nonsense questions such as "what colourless green ideas sleep furiously?" from language models. Furthermore, GenAI is an essential component of an emerging pseudoreligious meta-narrative of escape identified by Gebru & Torres (2024): "*What ideologies are driving the race to attempt to build AGI? [...] we trace this goal back to the Anglo-American eugenics movement, via transhumanism. [...] we delineate a genealogy of interconnected and overlapping ideologies that we dub the 'TESCREAL bundle,' where the acronym 'TESCREAL' denotes 'transhumanism, Extropianism, singularitarianism, (modern) cosmism, Rationalism, Effective Altruism, and longtermism'*".

Camus' existentialist view offers a non-escapist alternative that stares meaninglessness in the face and from it derives freedom. This freedom is both liberating and burdensome. We are at liberty to choose, but are also responsible for bearing the burden of the consequences. The lightness of being can thus be unbearable. It is through confronting this anxiety that individuals can make deliberate and meaningful choices, shaping their intentions, and by extension, their essence.

GenAI has implications for the intention of professional programmers and casual ones alike. The introduction poses the question "what is programming?", and we can now see a second reading of this question which asks not for a definition of an activity, but of an aspiration or identity. As GenAI solves the problem of control, of specifying behaviour, the aspiration shifts to intent. Intent precedes control. To be a programmer is therefore not to be one who specifies behaviour, but one who forms authentic, meaningful intentions for behaviour.

6. Speculative Scenarios for Intentional Programming

The optimism of the "intention is all you need" narrative does posit a legitimate observation concerning the behavioural economics of software production. GenAI makes the production of bespoke software vastly cheaper. One can view existentialism as a response to the loss of the "grand narratives" of modernity. But software has still been constrained by the grand narratives of capitalism and utility – until now. To write a program required *investment* of time and hard-earned expertise, exerting pressure on programs to be valuable, robust, and reusable. Where they did not place an outright barrier, the investment costs of programming disincentivised exploration, error, and disposal. Within this frame story hitherto sits

the universe of programmer psychology and behaviours: from authoring code to code comprehension, from knowledge sharing and documentation to debugging, from learning barriers to attention investment, from API design to autocomplete. Almost the entire diversity of experience of programmers, professional or casual, that our research community has so carefully documented and explained for the last half-century, has dwelt in the shadow of the market's invisible hand.

As the hand is withdrawn, one might ask how programmers can respond, in a microcosm of the existential dilemma, to the liberating yet burdensome freedom granted by GenAI. As far as practical advice (i.e., “implications for design[ing your life]”) is concerned, existentialists advise embracing one's freedom to shape life, living authentically, accepting the absurd, confronting anxiety, and seeking engagement with the world as ways to form meaningful intentions. What this might mean for programmers, and interaction with GenAI, can be sketched in a few speculative scenarios:

- **Intentional coding retreats:** The programmer steps away from her standard way of working to participate in an intentional coding retreat. Here, the programmer reconnects with the craft of coding without the assistance of AI tools. This allows the programmer to explore and reaffirm personal coding styles and problem-solving approaches. For example, a programmer accustomed to relying on AI for debugging might rediscover the satisfaction of manually untangling complex code, thus reaffirming their individual capability and creative freedom.
- **AI as muse:** AI suggests an unusual, contradictory, or incorrect algorithmic approach, which the programmer then refines and transforms with personal insights and expertise. The tool is not a crutch but a source of inspiration.
- **Programming with provocations:** programming environments include prompts or questions to stimulate deeper thinking about the purpose and potential impact of the code being written. This can help programmers reconnect with their motivations and aspirations.
- **Programming with constraints:** intentional constraints are introduced to programming projects, much like the practices of constrained writing.¹¹ Programmers already practice genres of constrained programming for pleasure, such as “code golf” (writing the shortest possible program with a certain behaviour) or “quines” (inputless programs that produce only their own source code as output). By deliberately limiting certain resources or imposing unique challenges, programmers can stimulate creativity and craft intentional solutions.
- **Deviation practice:** in the education of professional programmers, exercises are developed that require intentional deviation from established patterns. By practising the precise skill of breaking away from standard solutions, programmers may more readily acquire the conscious muscle and desire for forming unique intentions and exploring novel paths.
- **Intentionality metrics:** tools display metrics that evaluate the degree of human intention in the creative process (noting that these metrics are necessarily reductionist proxies and may become subject to Goodhart's/Campbell's law). For example, a generative design tool might analyse the uniqueness of user queries and the divergence of the output from standard templates. Visibilising the invisible effects of mechanised convergence may encourage users to engage more deeply with the work and make conscious, deliberate, individual choices.
- **Participatory AI artefacts:** artefacts are intentionally left incomplete by AI, requiring human participation¹² to finalise. For instance, a participatory tool generates the outline of a web design but leaves decisions about colour schemes and typography to the user. Conversely, a tool refuses to generate an outline, requiring the user to form a rough intention independently, before assisting by filling in details.

¹¹E.g., see discussion of conceptual writing in [Sarkar \(2023b\)](#).

¹²though not “collaboration” ([Sarkar, 2023a](#))

These speculations are not meant to be concrete proposals, but rather simply representative ideas of a future where the existentialist values of freedom, authenticity, and intentionality are preserved and enhanced through GenAI. They are limited in vision, representing only the lines of sight from where we stand today, and unable to anticipate the adjacent possibles of where we might travel.

7. Conclusion

Programming is undeniably changing under the influence of GenAI. Intention appears to be all one needs to create software. But the notion that GenAI offers a neutral, unencumbered path to realising intentions is a mirage. Contrary to the assumption that GenAI merely executes human intentions, it also shapes them. At the very least, GenAI can induce “mechanised convergence”, homogenising creative output, and reducing diversity in thought. There is therefore a risk of creating a “stuck” culture that recycles an old set of convergent ideas instead of fostering a new set of divergent ones.

In seeking a way through this problem we have encountered a variety of sources that we can draw upon to precipitate the active cultivation of intention: evolutionary pressures, the need for dignity and self-actualisation, constraints, subversion, materiality, and resistance. Finally, we discussed how the problem of intention resonates with the existentialist pursuits of freedom, identity, and authenticity. While this discussion of existentialism is necessarily cursory, limited, flawed, and provisional, its aim has been to situate the problems posed by GenAI to intentionality in the broadest possible scope.¹³

Programming must go beyond specification and embody the active cultivation of intentions. Existentialist philosophy offers a proactive, prescriptive framework for understanding the formation of human intentions as a process that ought to be held as deeply personal, ethically charged, and fundamentally free. It teaches us that to be human is to be involved in a continuous project of becoming. After all – one is not born, but rather becomes, a programmer.

8. Acknowledgements

Thanks to Sean Rintel and Lev Tankelevitch for helping review drafts of this paper. I am especially grateful to Ava Scott and Richard Banks for their generous and helpful reflections.

References

- Alemohammad, S., Casco-Rodriguez, J., Luzi, L., Humayun, A. I., Babaei, H., LeJeune, D., . . . Baraniuk, R. G. (2023). *Self-Consuming Generative Models Go MAD*. Retrieved from <https://arxiv.org/abs/2307.01850>
- Anderson, B. R., Shah, J. H., & Kreminski, M. (2024). Homogenization effects of large language models on human creative ideation. *arXiv preprint arXiv:2402.01536*.
- Arnold, K. C., Chauncey, K., & Gajos, K. Z. (2020). Predictive text encourages predictable writing. In *Proceedings of the 25th International Conference on Intelligent User Interfaces* (p. 128–138). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3377325.3377523> doi: 10.1145/3377325.3377523
- Arora, P. (2022). FemWork: Critical Pivot towards Design for Inclusive Labor Futures. In *2022 CHI Conference on Human Factors in Computing Systems (Closing Keynote)*.
- Basman, A. (2016). Building software is not a craft. *Proceedings of the Psychology of Programming Interest Group, 142*.
- Benjamin, R. (2024). *Imagination: A manifesto (a norton short)*. WW Norton & Company.
- Benjamin, W. (1935). The work of art in the age of mechanical reproduction, 1936. *New York*.
- Blackwell, A. F. (2002). What is programming? In *PPIG* (Vol. 14, pp. 204–218).

¹³Camus (1942) describes existence (suicide) as the only truly serious philosophical problem.

- Blackwell, A. F. (2020). Objective functions:(in) humanity and inequity in artificial intelligence. *Science in the ForeSt, Science in the PaSt*, 191.
- Bohacek, M., & Farid, H. (2023). *Nepotistically trained generative-ai models collapse*. Retrieved from <https://arxiv.org/abs/2311.12202>
- Buçınca, Z., Malaya, M. B., & Gajos, K. Z. (2021, apr). To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. *Proc. ACM Hum.-Comput. Interact.*, 5(CSCW1). Retrieved from <https://doi.org/10.1145/3449287> doi: 10.1145/3449287
- Cai, A., Arawjo, I., & Glassman, E. L. (2024). Antagonistic AI. *arXiv preprint arXiv:2402.07350*.
- Camus, A. (1942). *The myth of sisyphus* (J. O'Brien, Trans.). France: Éditions Gallimard (in French), Hamish Hamilton (in English).
- Carroll, J. M., & Rosson, M. B. (1987). Paradox of the active user. In *Interfacing thought: Cognitive aspects of human-computer interaction* (pp. 80–111).
- Carse, J. P. (1986). *Finite and infinite games*. New York, NY: Free Press.
- Chalmers, M., & MacColl, I. (2003). Seamless and seamful design in ubiquitous computing. In *Workshop at the crossroads: The interaction of HCI and systems issues in UbiComp* (Vol. 8).
- Cox, A. L., Gould, S. J., Cecchinato, M. E., Iacovides, I., & Renfree, I. (2016). Design frictions for mindful interactions: The case for microboundaries. In *Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems* (pp. 1389–1397).
- de Beauvoir, S. (1948). *The ethics of ambiguity* (B. Frechtman, Trans.). Citadel Press Publishing, A Subsidiary of Lyle Stuart Inc.
- Deleuze, G. (1992). Postscript on the societies of control. *October*, 59, 3–7. Retrieved 2024-06-04, from <http://www.jstor.org/stable/778828>
- Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., ... Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*(24-013).
- Dennett, D. C. (1971). Intentional systems. *The journal of philosophy*, 68(4), 87–106.
- Doshi, A. R., & Hauser, O. (2023, Aug). Generative artificial intelligence enhances creativity but reduces the diversity of novel content. (Available at SSRN: <https://ssrn.com/abstract=4535536> or <http://dx.doi.org/10.2139/ssrn.4535536>)
- Flanagan, M. (2009). *Critical play*. London, England: MIT Press.
- Foucault, M. (1977). *Discipline and punish*. New York, NY: Pantheon Books.
- Gaver, W. W., Beaver, J., & Benford, S. (2003). Ambiguity as a resource for design. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 233–240).
- Gebru, T., & Torres, É. P. (2024). The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. *First Monday*.
- Gerstgrasser, M., Schaeffer, R., Dey, A., Rafailov, R., Sleight, H., Hughes, J., ... Koyejo, S. (2024). *Is model collapse inevitable? breaking the curse of recursion by accumulating real and synthetic data*. Retrieved from <https://arxiv.org/abs/2404.01413>

- Gould, S. J. J., Brumby, D. P., & Cox, A. L. (2024). Chatldr – you really ought to check what the llm said on your behalf. In *Extended abstracts of the 2024 chi conference on human factors in computing systems*. New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3613905.3644062> doi: 10.1145/3613905.3644062
- Green, T. R. (1989). Cognitive dimensions of notations. *People and computers V*, 443–460.
- Heidegger, M. (1927). *Being and time* (J. Macquarrie & E. Robinson, Trans.). SCM Press.
- Hollanek, T. (2019). Non-user-friendly: Staging resistance with interpassive user experience design. *A Peer-Reviewed Journal About*, 8(1), 184–193.
- Hume, D. (1748). *An enquiry concerning human understanding*.
- Jacob, P. (2023). Intentionality. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy* (Spring 2023 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2023/entries/intentionality/>.
- Jakesch, M., Bhat, A., Buschek, D., Zalmanson, L., & Naaman, M. (2023). Co-Writing with Opinionated Language Models Affects Users' Views. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3544548.3581196> doi: 10.1145/3544548.3581196
- Jaspers, K., & Saner, H. (1932). *Philosophie* (Vol. 1). J. Springer Berlin.
- Kery, M. B., & Myers, B. A. (2017). Exploring exploratory programming. In *2017 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)* (pp. 25–29).
- Kierkegaard, S. (1843). *Fear and trembling*. Denmark: First authorship (Pseudonymous). (Published in English in 1919 – first translation)
- Kierkegaard, S. (1844). *The concept of anxiety* (R. Thomte, Trans.). Denmark. (Published in English in 1946)
- Kierkegaard, S. (1849). *The sickness unto death*.
- Lee, M., Blackwell, A., & Sarkar, A. (2024). Predictability of Identifier Naming with Copilot: A Case Study for Mixed-Initiative Programming Tools. *Proceedings of the 35th Annual Conference of the Psychology of Programming Interest Group (PPIG 2024)*.
- Liu, M. X., Sarkar, A., Negreanu, C., Zorn, B., Williams, J., Toronto, N., & Gordon, A. D. (2023). “What It Wants Me To Say”: Bridging the Abstraction Gap Between End-User Programmers and Code-Generating Large Language Models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3544548.3580817> doi: 10.1145/3544548.3580817
- Malafouris, L. (2019). Mind and material engagement. *Phenomenology and the cognitive sciences*, 18(1), 1–17.
- May, R. (1975). *The courage to create*. New York, NY: WW Norton.
- McLuhan, M. (1964). *Understanding media: The extensions of man*. McGraw-Hill. (First edition)
- Mele, A. (2019). Free will and neuroscience: decision times and the point of no return. In *Free will, causality, and neuroscience* (pp. 83–96). Brill.

- Mikolov, T., Yih, W.-t., & Zweig, G. (2013). Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 conference of the north american chapter of the association for computational linguistics: Human language technologies* (pp. 746–751).
- O’Neill, J. (1986). The disciplinary society: from weber to foucault. *British Journal of Sociology*, 42–60.
- Pfaller, R. (2017). *Interpassivity: The aesthetics of delegated enjoyment*. Edinburgh University Press.
- Postman, N. (1985). *Amusing ourselves to death*. Viking Books.
- Prather, J., Reeves, B. N., Denny, P., Becker, B. A., Leinonen, J., Luxton-Reilly, A., . . . Santos, E. A. (2023, nov). “it’s weird that it knows what i want”: Usability and interactions with copilot for novice programmers. *ACM Trans. Comput.-Hum. Interact.*, 31(1). Retrieved from <https://doi.org/10.1145/3617367> doi: 10.1145/3617367
- Rittel, H. W., & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy sciences*, 4(2), 155–169.
- Robinson, D., Cabrera, C., Gordon, A. D., Lawrence, N. D., & Mennen, L. (2024). Requirements are all you need: The final frontier for end-user software engineering. *arXiv preprint arXiv:2405.13708*.
- Rosen, B. M. (1979). *VISICALC: Breaking the Personal Computer Bottleneck*. <http://bricklin.com/history/rosenletter.htm>. (Accessed 06-08-2024)
- Rossmly, B., Terzimehić, N., Döring, T., Buschek, D., & Wiethoff, A. (2023). Point of no undo: Irreversible interactions as a design strategy. In *Proceedings of the 2023 chi conference on human factors in computing systems* (pp. 1–18).
- Rupert, R. D. (2004). Challenges to the hypothesis of extended cognition. *The Journal of philosophy*, 101(8), 389–428.
- Sarkar, A. (2016a). Constructivist Design for Interactive Machine Learning. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (p. 1467–1475). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/2851581.2892547> doi: 10.1145/2851581.2892547
- Sarkar, A. (2016b). *Interactive analytical modelling* (Tech. Rep. No. UCAM-CL-TR-920). University of Cambridge, Computer Laboratory. Retrieved from <https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-920.pdf> doi: 10.48456/tr-920
- Sarkar, A. (2023a). Enough With “Human-AI Collaboration”. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3544549.3582735> doi: 10.1145/3544549.3582735
- Sarkar, A. (2023b). Exploring Perspectives on the Impact of Artificial Intelligence on the Creativity of Knowledge Work: Beyond Mechanised Plagiarism and Stochastic Parrots. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3596671.3597650> doi: 10.1145/3596671.3597650
- Sarkar, A. (2023c). Should Computers Be Easy To Use? Questioning the Doctrine of Simplicity in User Interface Design. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3544549.3582741> doi: 10.1145/3544549.3582741

- Sarkar, A. (2023d). Will Code Remain a Relevant User Interface for End-User Programming with Generative AI Models? In *Proceedings of the 2023 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (p. 153–167). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3622758.3622882> doi: 10.1145/3622758.3622882
- Sarkar, A. (2024, sep). AI Should Challenge, Not Obey. *Communications of the ACM*. Retrieved from <https://doi.org/10.1145/3649404> (Online First) doi: 10.1145/3649404
- Sarkar, A., Gordon, A. D., Negreanu, C., Poelitz, C., Srinivasa Ragavan, S., & Zorn, B. (2022, September). What is it like to program with artificial intelligence? In *Proceedings of the 33rd Annual Conference of the Psychology of Programming Interest Group (PPIG 2022)*.
- Sarkar, A., Xu, X. T., Toronto, N., Drosos, I., & Poelitz, C. (2024). When Copilot Becomes Autopilot: Generative AI's Critical Risk to Knowledge Work and a Critical Solution. In *EuSprIG Proceedings*.
- Sartre, J.-P. (1943). *Being and nothingness* (H. E. B. (1st English translation) & S. R. (2nd English translation), Trans.). France: Éditions Gallimard, Philosophical Library. (Published in English in 1956)
- Sharma, N., Liao, Q. V., & Xiao, Z. (2024). Generative Echo Chamber? Effect of LLM-Powered Search Systems on Diverse Information Seeking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3613904.3642459> doi: 10.1145/3613904.3642459
- Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). Ai models collapse when trained on recursively generated data. *Nature*, *631*(8022), 755–759.
- Snyder, L. (2012). *The philosophical breakfast club*. New York, NY: Broadway Books.
- Stokes, P. D. (2005). *Creativity from constraints: The psychology of breakthrough*. Springer Publishing Company.
- Tankelevitch, L., Kewenig, V., Simkute, A., Scott, A. E., Sarkar, A., Sellen, A., & Rintel, S. (2024). The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3613904.3642902> doi: 10.1145/3613904.3642902
- Turner, P. (2016). Distributed, external and extended cognition. *HCI Redux: The Promise of Post-Cognitive Interaction*, 75–98.
- Welsh, M. (2022, dec). The end of programming. *Commun. ACM*, *66*(1), 34–35. Retrieved from <https://doi.org/10.1145/3570220> doi: 10.1145/3570220
- Winner, L. (1980). Do artifacts have politics? *Daedalus*, 121–136.
- Zamfirescu-Pereira, J., Wong, R. Y., Hartmann, B., & Yang, Q. (2023). Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3544548.3581388> doi: 10.1145/3544548.3581388