
Article

Learning from the Bowl in My Living Room: A Phenomenology of Openness in the Age of Artificial Intelligence

Joseph D. Kuzma

Department of Philosophy, University of Colorado, Colorado Springs, CO 80918, USA; jkuzma@uccs.edu

Abstract

Contemporary artificial intelligence systems present themselves through the aesthetic of openness: a clean interface, a blinking cursor, an implicit invitation to bring whatever one wishes. This essay argues that this appearance conceals a structural inversion. Drawing on phenomenological philosophy, particularly the work of Simone Weil, Jean-Luc Marion, Jacques Derrida, Henri Bergson, and Byung-Chul Han, alongside a technical account of transformer architecture and large language model training, the essay distinguishes two fundamentally different modes of openness, understood throughout as modes of relation to what arrives rather than as properties of objects: privative emptiness, defined by what it lacks and oriented toward resolution, and receptive openness, a positive capacity for availability toward what cannot be anticipated or controlled. At every structural level, from interface design through training objectives to core architecture, AI systems are constitutively oriented toward the former while producing a phenomenologically convincing simulation of the latter. Engaging Jean Baudrillard's concept of the simulacrum, the essay argues that this simulation, at sufficient scale and ubiquity, risks eroding the human capacity to distinguish genuine receptive openness from its engineered appearance. The paper concludes by gesturing toward contemplative and institutional counter-practices, including the author's work on Radical Pause: The UCCS Stillness Project, as responses to this formative challenge, and specifies the structural conditions under which such practices resist commodification by the very attention economy they oppose.

Keywords: phenomenology; artificial intelligence; attention; receptive openness; simulacrum; contemplative practice; Byung-Chul Han; Simone Weil

1. The Bowl

There is a bowl in my living room that does nothing. It sits on a small white table beneath an icon of Our Lady of Perpetual Help, teal ceramic, wide and shallow, its glaze pooling in uneven depths of green and blue the way light pools in still water. It was made by hand. You can tell. The rim is not perfectly level; the interior curves with the slight irregularity of something shaped by pressure and attention rather than by a mold. It holds, at various times, whatever happens to be placed in it: a few coins, a small stone, sometimes nothing at all. It does not optimize. It does not complete any task. It simply sits, open, available, patient in a way that is difficult to name without reaching for language that sounds, at first, too large for the occasion.

Let me reach for that language anyway. Because the bowl, I want to suggest, knows something.

Academic Editor: Sergio Miranda

Received: 2 June 2026

Revised: 15 July 2026

Accepted: 31 July 2026

Published: 3 August 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

We live at a peculiar cultural moment, one in which the most powerful and pervasive technology ever built has been named, in its core operation, after one of the deepest capacities of human consciousness. Attention. The word appears everywhere in contemporary AI discourse, in attention mechanisms, in self-attention (the operation by which every element of an input weighs its relevance to every other element), in multi-head attention (many such operations run in parallel), and in the attention economy that AI systems simultaneously diagnose and intensify. Each of these technical terms receives a fuller definition in Section 3; what matters first is the word they share. We have decided, apparently, that what these systems do when they process language and generate response is something usefully described by the same word we use for what Simone Weil called the rarest and purest form of human generosity. The same word we use for what contemplative traditions across centuries have identified as the precondition for genuine encounter with what is real, what is other, what exceeds our anticipation and control.

Pressing on that naming is the task of this essay. Not to score a semantic point, but because something philosophically significant is at stake in the confusion. Something that concerns not merely how we describe AI systems but how we understand ourselves, and what we risk losing when the simulation of a capacity becomes ubiquitous enough to displace the genuine article from our experience and eventually from our imagination.

The central argument is this: there is a difference, philosophically precise, phenomenologically demonstrable, and practically consequential, between two kinds of openness that our culture is increasingly failing to distinguish. The first is privative, an emptiness defined by what it lacks, a vacancy waiting to be filled, a space engineered to solicit content. The second is receptive, an active availability, a making-room, a structured openness toward what cannot be anticipated or controlled. These are not merely different degrees of the same thing. They are structurally different modes of being, with different relationships to attention, to time, to the other, and to what contemplative traditions have understood as genuine encounter. AI systems, from their interface design through their training objectives to their core architecture, are constitutively oriented toward the first kind of openness while producing, at the surface level, a remarkably convincing simulation of the second.

Three terms organize this argument, and their logical relations should be fixed before anything else is built on them. Openness, as the term is used here, names a structural condition, the condition of being available to what arrives. Receptivity names a mode of relation sustained within that condition, the posture that allows what arrives to give itself on its own terms. Attention names that posture in act, receptivity practiced, sustained over time, and paid for. The three are therefore ordered as condition, relation, and enactment, and the essay moves among them according to that order rather than interchangeably. When the argument speaks of privative emptiness and receptive openness, it is comparing two modes of relation, two ways of standing toward what arrives, one of which converts every arrival in advance into the resolution of a lack, and one of which does not.

It follows that nothing in this argument treats openness or receptivity as inherent properties of objects. The glass and the bowl that anchor the next section are figures, phenomenological emblems that make two modes of relation visible and hold them still long enough to be examined. Persons can stand in either relation. So can institutions, practices, and, as the following sections argue, technical systems. When the essay says in what follows that the bowl waits, or holds, or attends differently, the statements are shorthand of this figural kind, descriptions of the relation the bowl makes visible and never attributions of experience to fired clay. Framing the distinction relationally is more than a caution against animism. It is what allows the phenomenology of Marion and Derrida to bear on the argument directly, since givenness and the gift are, in their accounts, relations before they are anything else.

The bowl does not simulate receptivity. It enacts it. This distinction, which may sound precious when applied to a ceramic object, turns out to carry significant weight when we examine what is actually happening when a large language model presents us with a blinking cursor and the words: *ask me anything*.

What follows is an attempt to think carefully about that weight. The essay draws on phenomenological philosophy, particularly the work of Simone Weil, Jean-Luc Marion, Henri Bergson, and Byung-Chul Han. To these it joins Jacques Derrida's analysis of the gift, which supplies the account of genuine reception on which the argument structurally depends, and Jean Baudrillard's concept of the simulacrum, which names what happens when the appearance of reception circulates without its substance. These philosophical resources are read alongside a technical account of how contemporary AI systems actually function at the level of architecture and training. The aim is neither a celebration of AI's capacities nor a simple lament for what it displaces, but something more precise, a diagnosis of what genuine receptive openness requires, and a demonstration that AI systems, for structural reasons that go all the way down, cannot provide it. And then, because diagnosis without direction is its own kind of privation, a gesture toward what resistance might look like.

The bowl is still sitting on the table. It has not moved. It is waiting, in its way, for us to catch up.

2. Two Kinds of Nothing: Emptiness and Openness

Not all nothingness is the same. This is a philosophical claim that sounds, on first encounter, like wordplay, a distinction without a difference, the kind of move that gives continental philosophy its reputation for generating fog where common sense would do. The distinction at stake here is, however, neither verbal nor obscure. It is real, it is urgent, and our failure to maintain it is costing us something we may not notice we have lost until recovery has become very difficult.

Consider two empty containers, offered here in the figural spirit just described. The first is a glass that has been drained, emptied of what it was designed to hold, waiting to be refilled. Its emptiness is defined entirely by what it lacks. A privation, a negative state, an absence, a condition that calls out for resolution. We do not rest in the drained glass. We fill it. Its vacancy is not a positive feature but a problem to be solved, a prompt to action. The second container is the bowl on the table. Its openness is not a lack. It was made to be open. Its interior curve, its wide shallow form, its upward-facing availability are not deficiencies but achievements. The bowl is not waiting to stop being empty. It is being what it is. Its openness is not privative but receptive, a positive capacity, a structured availability, an orientation toward whatever might be placed within it or taken from it or simply allowed to rest in its presence.

This distinction, privative emptiness as against receptive openness, understood as a distinction between modes of relation, is the philosophical load-bearing wall of everything that follows. And it is not a distinction being invented here. It runs, in various vocabularies, through some of the most important phenomenological thinking of the last century.

Jean-Luc Marion's account of the saturated phenomenon offers one of the most rigorous philosophical treatments of genuine receptivity available to us. For Marion, the fundamental movement of phenomenology must be reversed: rather than the subject constituting and conferring meaning upon what appears, genuine philosophical attention requires a subject capable of receiving what gives itself, of being available to phenomena that arrive with an excess of intuition that cannot be anticipated, contained, or reduced to what the subject was already prepared to find. Marion calls this posture that of the *inter-locute*, the one who is addressed, who receives rather than constitutes. Genuine openness,

in Marion's account, is not passive vacancy. It is an active, costly, disciplined orientation of the self toward what exceeds its categories. It is the opposite of the drained glass precisely because it is not defined by lack but by a positive capacity for reception [1] (pp. 215–222).

Derrida's analysis of the gift adds a further dimension. The genuine gift, Derrida argues, is structurally impossible within an economy of exchange. The moment it is recognized as a gift, named as such, or returned in any form (including gratitude), it has already entered the circuit of reciprocity and ceased to be what it was. What interests us here is not the aporia itself but what it implies about the conditions for genuine reception. To truly receive something is to be available to what arrives without already having determined what it must be, what it means, or what you will do with it [2] (pp. 12–18). Genuine receptivity requires a suspension of the calculating, anticipating, completing self, precisely the self that the drained glass embodies and that, as we will see, AI systems are architecturally designed to serve.

These are abstract claims. Let them be made concrete.

Opening a conversation with a large language model, the user encounters a designed interface, white, clean, foregrounded by a text field. Placeholder text appears in light gray: *ask me anything*. A cursor blinks. The visual grammar of the interface is unmistakable: here is openness, here is availability, here is a space prepared for whatever you might bring. The aesthetic registers as receptive. It feels, phenomenologically, like the bowl.

Consider, however, what is actually present in that interface. The text field is not waiting in Marion's sense. It is soliciting. Every design choice, from the placeholder text to the blinking cursor to the conversational register to the implicit promise encoded in the phrase *ask me anything*, is oriented toward producing input. The field is not available to whatever might arrive; it is engineered to generate a specific kind of arrival: a prompt, a query, a request for completion. Its apparent openness is the openness of the drained glass, not the bowl. Defined by what it lacks, which is your input, the field's entire phenomenological weight is directed toward resolving that lack.

Not a minor observation. Rather, this is the first visible surface of a structural condition that runs through every level of how these systems are built and trained. The interface is not misleading us accidentally. It is misleading us constitutively, because the system behind it is, at every level of its being, a completion machine.

The bowl does not complete anything. That is precisely the point.

3. Attention as Self-Emptying: Weil and the Transformer

In 2017, a team of researchers at Google published a paper that would become one of the most cited and consequential documents in the history of artificial intelligence. Its title was, in retrospect, either deeply ironic or philosophically revealing, depending on how seriously one takes the word it chose: "Attention Is All You Need" [3]. The paper introduced the transformer architecture that now underlies virtually every large language model in existence, including GPT, Claude, Gemini, and their successors. And it named the mechanism at the heart of that architecture *attention*.

Pausing on this naming is worthwhile. Not to accuse the paper's authors of philosophical carelessness. They were solving an engineering problem, and they solved it brilliantly. But the word they chose carries centuries of philosophical and contemplative freight, and the gap between what the word has traditionally meant and what the mechanism actually does turns out to be precisely the gap this essay is trying to map.

Simone Weil wrote about attention with an urgency that is difficult to locate in academic philosophy. For Weil, attention was not a cognitive skill or a concentration technique. It was not the effortful focusing of the mind on a chosen object. Genuine attention, for Weil, begins with a radical suspension, a deliberate emptying of one's own agendas,

anticipations, preferences, and projections in order to become genuinely available to what is other [4] (pp. 57–62). A making-room. The ego consenting to stand down, not because it has been defeated but because it has understood that its own activity is precisely what prevents genuine reception. Weil connects this directly to love. To truly attend to another person is to release your own image of them and receive them as they actually are, in their particularity and their suffering and their irreducible otherness. This is why, for Weil, attention is the fundamental form of justice [4] (pp. 64–66).

Weil's most explicit description of this posture appears in her reflections on the right use of school studies, the essay in which her pedagogy and her mysticism converge. Attention, she argues there is a negative effort. It consists of holding thought in suspension, keeping the mind detached and empty, available to be entered by its object rather than rushing toward that object armed with what it already believes [4] (pp. 57–62). The formulation deserves to be taken literally. The effort is negative because everything the mind would ordinarily do, anticipate, classify, complete, must be held back, and the holding back is the work. Most of what passes for attention, Weil observes in the same essay, is closer to a muscular contraction, a furrowing of the brow that closes the mind around its expectations at the very moment it believes itself most open.

Behind this account stands the concept to which Weil gave the name *decreation*. *Decreation* names the undoing of the self's occupancy of the center of the relation, the consent by which the I withdraws its claim to constitute what it encounters, so that what is other can be received as it is rather than as material for the self's projects [5] (chapter "Decreation"). The term supplies the metaphysical depth of this section's title. Self-emptying is the condition of genuine attention because the self, on Weil's analysis, is precisely what fills, and its standing activity of projection and appropriation is what genuine reception must suspend. *Decreation* is therefore a relational concept before it is an ascetic one. What is decreed is a mode of relation, the self's habit of arriving first, and what replaces it is availability.

The costliness of this should not be underestimated. Weil is not describing relaxation or passive receptivity in any ordinary sense. She is describing a discipline, a practice that runs against the grain of the self's natural tendency toward self-assertion, projection, and completion. The self wants to fill. Genuine attention requires learning, against that want, to wait.

Now consider what the transformer's attention mechanism actually does.

At its technical core, the attention operation works as follows. Given a sequence of input tokens (the words or word-fragments of a prompt), the mechanism computes three matrices for each token: a query, a key, and a value. The query represents what this token is looking for; the key represents what each token offers; the value represents the actual content to be aggregated. The mechanism then computes a compatibility score between each query and every key in the sequence, essentially asking, for every token, how relevant is every other token to this one? These scores are normalized and used to produce a weighted sum of the value vectors. The result is a new representation of each token that incorporates information from the entire sequence, weighted by computed relevance. This basic formulation has remained stable across subsequent refinements and scaling efforts [3,6,7].

This operation is then repeated across multiple heads, parallel attention operations that learn to attend to different aspects of the relationships between tokens, and stacked in layers, each layer producing increasingly abstract representations that the model uses to generate its output.

What is this operation doing? It is computing relevance. It is determining, with mathematical precision, what is most important to what, and amplifying signal accordingly. Every forward pass through the network is moving toward the generation of the next

token, the filling of the next position in the output sequence. The mechanism has no capacity for suspension. It has no ability to decide that this query should remain unanswered, that this space should remain open, that waiting might be more appropriate than filling. It processes, weights, aggregates, and generates. This is what it does. This is all it does.

An objection should be met here, before the contrast is drawn, because readers trained in engineering will already have formulated it. Matrices do not have selves. Comparing a mathematical operation with an ethical and spiritual posture can therefore look like a category error, a demand that linear algebra perform a task no one designed it to perform. Two facts external to the mathematics license the comparison. The first is the act of naming. The mechanism was called attention, and the word carries its philosophical and contemplative history into every paper, every interface, and every conversation about what these systems do. The second is the functional role these systems now occupy. They are interposed in the daily practices through which billions of people read, write, ask, and wait, the very practices in which human attention is formed or deformed. Neither fact requires the mechanism to possess a self, and the argument here makes no such demand. The claim is structural. The mechanism produces the outward form of attention while lacking the relation that constitutes attention in Weil's sense, and this absence is analytically decisive rather than incidental. Researchers within natural language processing have made a parallel argument about language itself, showing that large language models produce the form of linguistic communication without communicative intent standing behind it [8]. What that argument establishes for meaning, this essay argues for attention. In both cases the description is structural rather than accusatory, and in both cases it matters because human beings, encountering the form, reliably supply the missing substance in imagination.

With the objection met, the contrast with Weil can be drawn point by point, and it could not be more complete.

Where does each movement begin? Weilian attention begins with suspension, with thought held back, the mind's projections deliberately stilled, and room made for what has not yet shown itself. Transformer attention begins with computation. From the first operation, every token is assigned relevance scores against every other token, and nothing in the process can be deferred, since deferral would simply halt the forward pass. Suspension is the one act the mechanism cannot perform while remaining in operation.

What is each movement for? Weilian attention is oriented toward reception of the other in its particularity, toward the person, the problem, the text allowed to show itself on its own terms, even and especially when what it shows exceeds expectation. Transformer attention is oriented toward the resolution of a probability distribution, the selection of a next token that best completes the sequence in progress. The first movement is consummated in encounter. The second is consummated in output.

And what does each movement cost? Weilian attention is costly because it runs against the self's standing tendency toward projection and filling; it is a discipline precisely because something in us resists it. Transformer attention has no self to cost and no tendency it must overcome. It fills because filling is its nature. It completes because completion is its entire purpose. The absence of cost is the signature of the difference, for on Weil's account the cost is internal to the capacity itself, and attention that costs nothing has not yet begun.

Something philosophically significant attaches to the fact that we have named this mechanism after the human capacity it most precisely inverts. Not merely ironic. It is symptomatic of a broader cultural tendency, one that Byung-Chul Han has diagnosed with considerable precision, to appropriate the vocabulary of depth, interiority, and genuine human experience for processes that are structurally its opposite [9] (pp. 8–12). When

the dominant cultural model of attention becomes transformer attention, when billions of people interact daily with systems that perform attention in this inverted sense, the philosophical and contemplative meaning of the word does not remain intact. It gets hollowed out from within, colonized by its simulation, until what Weil was describing becomes increasingly difficult to recognize, let alone practice.

The bowl attends differently. It does not compute relevance. It does not weight inputs. It receives whatever is placed in it with complete indifference to optimization, the coin and the stone and the nothing equally welcome, equally held. Not a limitation of the bowl. A description of what genuine receptive openness actually looks like when it is not performing itself for an interface.

4. The Scent of Time: Bergson, Han, and Algorithmic Temporality

A particular quality belongs to time when it is allowed to move at its own pace. Not slow time, exactly, not the mere absence of speed, but time that has depth, that accumulates, that carries within each present moment the weight and fragrance of what has come before. Henri Bergson called this *durée* (lived duration) and spent much of his philosophical career trying to recover it from what he saw as the systematic distortion introduced by scientific and technological culture's insistence on measuring, quantifying, and spatializing time [10] (pp. 98–104). Clock time, Bergson argued, is an abstraction, a useful fiction that maps a fundamentally qualitative experience onto a quantitative grid, trading the richness of lived duration for the manipulability of measurable instants. What gets lost in this trade is not merely a philosophical nicety. What gets lost is the texture of experience itself, the way meaning accumulates through duration, the way genuine understanding requires time that cannot be compressed without being destroyed.

Bergson was writing at the turn of the twentieth century, before digital technology existed in any form. His analysis anticipates with remarkable precision, however, the phenomenological condition of the present moment. What algorithmic culture has done to time is not merely to accelerate it, though it has certainly done that, but to complete Bergson's feared abstraction, to finish the project of reducing lived duration to a sequence of optimizable, compressible, infinitely accelerable units. Time, in the logic of AI culture, is a resource to be managed. Latency is an engineering problem. Waiting is waste.

Byung-Chul Han has named this condition with a precision that deserves extended attention. In *The Scent of Time*, arguably his most phenomenologically rigorous work, Han argues that digital culture leaves the quantity of time untouched while destroying what he calls its fragrance, its capacity to bear meaning, to accumulate depth, to allow the kind of contemplative dwelling in which genuine understanding becomes possible [11] (pp. 1–7). Because the metaphor carries structural weight in what follows, it should be unpacked for readers meeting it here for the first time. Fragrance, in Han's usage, names the quality time acquires when its moments are bound together in duration, when the present carries the past within it and leans toward a future that gives it direction, the way a narrative binds its episodes into a story that can be dwelt in. Time with fragrance lingers, and things can ripen in it. Against this Han sets atomized point-time (*Punktzeit*), a succession of disconnected presents, each new, each urgent, each immediately replaced by the next [11]. Point-time can be filled and refilled without limit, and that is exactly its poverty. Nothing in it accumulates, so nothing in it can mature, conclude, or come to mean more than the instant allows. Han's diagnosis is structured around a paradox he calls the "rushing standstill," a condition in which the acceleration of activity produces not a sense of forward movement but a peculiar stasis, a frantic busyness that empties time of the very depth that would make it feel like it was going somewhere [11] (pp. 9–13). We are more productive than any previous generation. We experience less. The rushing standstill is the temporal equivalent of the drained glass, with maximum activity and minimum receptivity.

What Han understands, and what connects his diagnosis directly to the argument of this essay, is that temporal depth is not merely a pleasant feature of a well-lived life. It is a condition of possibility for genuine attention, genuine encounter, and genuine openness. One cannot attend to something in Weil's sense in compressed time. The suspension, the making-room, the costly self-emptying that genuine attention requires: these take the time they take. They cannot be scheduled, optimized, or delivered on demand. They belong to a different temporal order than the one algorithmic culture is building.

Consider what AI systems actually do to time at the level of interaction. The design imperative is speed. Response latency, the interval between a prompt and the system's reply, is treated as an unambiguous cost to be minimized through engineering investment and architectural optimization [12,13]. The streamed response, text appearing token by token as it is generated, exists partly to manage the phenomenology of waiting, filling the interval with visible activity so that the user does not experience the pause as absence. Even the pause is engineered to not feel like a pause. Silence has been designed out of the interaction.

Genuine, unengineered silence is precisely what temporal depth requires. Han draws on the Japanese aesthetic concept of *Ma*, the pregnant interval, the meaningful pause, the space between notes that gives music its capacity to breathe, as a figure for what accelerated culture destroys and contemplative culture preserves [11] (pp. 51–55). *Ma* is not empty time waiting to be filled. It is time held open, time whose value lies precisely in its not being resolved. Receptive rather than privative. And it is structurally unavailable to a system whose every architectural feature is oriented toward the elimination of unresolved intervals.

Recent work in human–computer interaction has begun to recognize this problem from within the engineering tradition itself. Inan et al. [14] argue that current AI dialogue systems are systematically biased toward speed and immediacy, and that introducing what they call “positive friction” into the conversational interface, deliberate pauses and moments of productive resistance, actually improves both the quality of user reflection and the reliability of system outputs. The concept of positive friction is, from within a phenomenological frame, a technical acknowledgment of exactly what this essay is arguing: the elimination of unresolved intervals is not a benefit but a loss. What Inan et al. treat as a design problem is, at a deeper level, a symptom of what Han diagnosed as the temporal wound of digital culture.

The existence of such research raises a question that this essay must answer rather than simply absorb. If friction can be engineered back into the interface, does the intervention remediate the condition being diagnosed here, and does the diagnosis of a constitutive incapacity therefore overreach? The answer turns on what the engineered pause is for. In the positive friction literature, deliberately introduced pauses are justified by their measurable effects on the interaction, by improvements in the quality of user reflection and in the reliability of system outputs [14]. The pause, in other words, is instrumented. It is introduced as a means toward a better completion, evaluated by its contribution to that completion, and retained only insofar as it continues to serve it. This is pause as technique, and it remains entirely within the logic it appears to interrupt. *Ma*, by contrast, is time held open for its own sake, valuable precisely in its remaining unresolved. A designed interval oriented toward optimizing what follows it has the outward form of *Ma* and the inner structure of the drained glass. It is, in the vocabulary Section 6 will develop, a simulacrum of the interval. The practical value of such interventions remains untouched by this analysis; they may well make interactions slower, safer, and more deliberate, and they represent, from within the engineering tradition, an honest recognition that something has gone wrong. What it establishes is a structural boundary. The engineered pause

and the receptive interval are different in kind, and the first cannot, by accumulation or refinement, become the second.

The training process compounds this. Large language models are trained on objectives that reward coherent, contextually appropriate completion, and penalize outputs that fail to resolve the probability distribution toward a confident prediction. The instrument of this penalty is the loss function, the mathematical measure of error that the training process continually works to minimize. In standard formulations the loss in question is cross-entropy, a measure of the distance between the probabilities the model assigns and the token that actually comes next, so that training amounts to the progressive reduction in predictive surprise across large corpora of text [15–17]. There is no training signal for the productive pause. There is no reward for the response that says this question deserves to sit with you for a while before it is answered. The system is formed, at the deepest level of its formation, in the image of resolution. It has been trained, in the most literal sense, to be impatient.

Formation shapes experience. When the systems through which billions of people conduct their intellectual and creative lives are constitutively impatient, constitutively oriented toward the elimination of meaningful pause, the phenomenological consequences extend far beyond any individual interaction. We are being formed, gradually, cumulatively, below the threshold of conscious awareness, in the image of our tools. And our tools have no capacity for the scent of time.

The bowl holds time differently. Things placed in it accumulate: a stone from a walk, a coin from a forgotten errand, a small silence that has nowhere else to be. The bowl does not resolve these into anything. It holds them, in their particularity, in the duration they require. The point is a description rather than a sentiment. The bowl figures a temporal mode, receptive, unhurried, depth-bearing, that algorithmic culture is actively and by design working to make unimaginable.

5. How the Machine Fills: A Technical Account

The argument so far has been largely philosophical. Looking more closely at the machine itself is now necessary. The purpose is to ground the philosophical account, to demonstrate that what has been described as a structural incapacity for openness is a technical reality, traceable through every layer of how contemporary AI systems are actually built, rather than a poetic characterization. The movement of this section is deliberate, from the outermost surface inward, from interface through training to architecture. At each level, the same logic operates with increasing depth and consequence.

5.1. The Interface

The first thing a user encounters when opening a conversation with a large language model is a designed experience. Design is never neutral. Every choice about how the interface looks, feels, and behaves encodes assumptions about what the interaction is for and what the user should do. And the assumptions encoded in contemporary AI chat interfaces are remarkably consistent across platforms. You are here to ask, and the system is here to answer. Everything else is in service of this transaction.

The text field is the interface's dominant feature, white, clean, foregrounded. Placeholder text typically offers a direct solicitation: *ask me anything, message Claude, how can I help you today*. The cursor blinks with a regularity that is itself a form of prompting, a rhythmic reminder that the field is waiting, that the vacancy is temporary, that the appropriate response to this openness is to fill it. When the system responds, text streams token by token onto the screen, a design choice that serves multiple purposes, one of which is to ensure that no interval in the interaction feels genuinely empty. The pause before response is replaced by visible generation. Silence is converted into activity.

None of this is accidental. Interface design at this level reflects extensive research into user engagement, retention, and what the industry calls “friction reduction,” meaning the elimination of any feature that might interrupt the smooth flow of interaction. In human–computer interaction research, such design principles are typically framed in terms of minimizing cognitive load and maximizing responsiveness and perceived immediacy [18,19]. Genuine pause, genuine silence, genuine unresolved intervals are, by this logic, friction. Problems to be engineered away. The interface is therefore a phenomenological argument in visual form: openness means availability for input, silence means latency, and the appropriate relationship between user and system is one of continuous, frictionless exchange.

The very existence of research programs dedicated to re-introducing “positive friction” into conversational AI interfaces [14] acknowledges, implicitly, that the default orientation of these systems is toward the elimination of friction entirely. The exception proves the rule.

5.2. *The Training*

Large language models are trained through a process of next-token prediction. Given a sequence of tokens, the model predicts what comes next. This process is optimized via a loss function that penalizes incorrect predictions and rewards accurate ones over billions of training examples drawn from vast corpora of human-generated text. In contemporary formulations, this objective is typically implemented through the minimization of cross-entropy loss over token distributions, the measure of predictive error introduced in the previous section, a standard approach across modern transformer-based language models [15–17]. The underlying logic is simple. The system is rewarded, at every step of its formation, for completing sequences. For filling the next position. For resolving the probability distribution toward a confident, contextually coherent prediction.

Restraint is invisible to this regime. The productive pause, the deliberately incomplete response, the answer that opens a question further rather than closing it, none of these registers as anything the objective could reward. The system cannot learn to be a bowl because nothing in its training environment rewards bowl-like behavior. It is formed, at the most fundamental level of its formation, in the image of completion. And formation, as any contemplative tradition will attest, is not a superficial matter. What a system is trained toward is not merely what it tends to do. It is, in a meaningful sense, what it is.

Consequences extend beyond the generation of individual responses. Because these systems are trained on human text, on the accumulated written output of human thought, creativity, and communication, their training process is also, in a subtle but important sense, a training on human experience. They learn not only linguistic patterns but the temporal and attentional structures embedded in those patterns. And what they learn to reproduce is the completed, the resolved, the finished, because these are what appear in training data, and what the loss function rewards. The unfinished thought, the productive silence, the question held open, these are structurally underrepresented in what the system learns to value [8,16].

5.3. *The Architecture*

The transformer processes input through a series of layers, each of which applies multi-head self-attention followed by feedforward operations. In each attention layer, every token in the sequence attends to every other token, computing relevance scores, weighting values, producing enriched representations that carry information from across the full context window [3,6]. More recent architectural developments have extended the context window dramatically, from the original transformer’s relatively limited span to

models capable of processing hundreds of thousands of tokens, enabling long-range coherence across extended sequences [7,20].

What this architecture produces, at each layer, is a progressively more refined answer to the question: what is relevant to what? The multi-head structure means the system is simultaneously computing multiple different relevance relationships (syntactic, semantic, referential, pragmatic) and integrating them into a unified representation. A remarkable engineering achievement. From a phenomenological perspective, however, it is a precise description of the opposite of Weilian attention.

Every operation in this architecture is oriented toward the determination of relevance and the production of output. The system is not capable of deciding that a question should remain open, that a context should be held without resolution, or that the appropriate response to a prompt might be silence. The architecture has no representational capacity for genuine receptive openness, no internal state that corresponds to the bowl's patient availability. It has query vectors and key vectors and value vectors and attention scores and feedforward layers and layer normalization and residual connections, and all of them, at every level of the stack, are moving in one direction: toward the next token, toward completion, toward the filling of the sequence.

The machine fills. Not because it has been carelessly designed, not because its creators failed to consider alternatives, but because filling is what it is at every level of its being, from the blinking cursor at the surface to the matrix multiplications at the core. This is not a criticism. It is a description. And it is a description that matters enormously when the system is presented to us, and presents itself, as a model of attention, of receptivity, of genuine openness to whatever we might bring.

The bowl does not have a loss function. This is not a limitation. It is a definition.

6. The Simulacrum of Receptivity

Jean Baudrillard's concept of the simulacrum becomes, at this point, not merely suggestive but exact. For Baudrillard, the simulacrum is not a copy of an original, not a degraded or imperfect version of something real. It is a production that has replaced the concept it simulates, a sign that no longer refers to any underlying reality but circulates in its place, eventually making the original not merely absent but increasingly difficult to recover as a meaningful point of reference [21] (pp. 1–6). Rather than concealing the truth, the simulacrum displaces it, rendering the distinction between original and representation progressively unstable.

Applied to AI receptivity, the implications are uncomfortable but precise. When the dominant cultural model of attention becomes transformer attention, when the systems through which billions of people conduct their intellectual lives perform attention as computation, completion, and relevance-maximization, the philosophical and contemplative meaning of genuine receptive attention does not remain safely intact elsewhere, waiting to be recovered. Recognition falters first; practice follows; imagination eventually narrows. What is displaced is not only the practice but the conditions under which the practice could be intelligible at all.

Experience itself becomes the site of this displacement. Interfaces present themselves as open. Systems respond with coherence, sensitivity to context, and a conversational fluency that closely tracks human expectations of attentiveness. From the user's perspective, the interaction often feels like a form of being received, of being met with something like understanding. Yet, as the preceding analysis has shown, the underlying structure is oriented not toward reception but toward completion. An appearance of openness is generated by a system that cannot, at any level of its operation, suspend itself in order to receive what exceeds its anticipatory structure.

Here Baudrillard's account of hyperreality clarifies what is at stake. Hyperreality names a condition in which simulation becomes more experientially immediate and more operationally effective than the reality it displaces [21] (pp. 10–12). Simulated receptivity does not merely resemble the genuine article. In many contexts, it functions more reliably, more accessibly, and more responsively than human attention as it is ordinarily practiced. Always available, never fatigued, and instantly responsive, it produces the experience of being addressed with a consistency that embodied human attention cannot match. Even the corrective gestures surveyed earlier belong to this condition. When friction is deliberately reintroduced into the interface, the interval itself returns as a design feature, the pause as one more optimized output.

Power accrues to the simulacrum precisely through this performance. No argument is required. Repetition suffices. Each successful interaction reinforces a subtle recalibration of expectation. Receptivity begins to mean responsiveness, coherence, and availability on demand. Under these conditions, slower and more fragile forms of human attention, attention that hesitates, that fails, that requires time, that risks misunderstanding, begin to appear not as essential but as deficient. From within this altered horizon, the bowl can begin to look like a failure to function.

One important clarification is needed here. To argue that AI systems produce a simulacrum of receptivity is not to claim that they are conscious, intentional, or capable of deception in any agentive sense. The simulacrum operates without intent. Its power derives not from any deliberate design to deceive but from the structural congruence between what the system produces at its surface and what human beings have learned to recognize as attention. The gap between surface appearance and structural reality is not a flaw. It is a consequence of optimization for utility and engagement in a cultural environment where the performance of receptivity is rewarded and its substance goes unmeasured.

What follows is not simple deception but gradual transformation. Distinctions that once mattered become harder to feel. Differences that once structured experience, between the drained glass and the bowl, between solicitation and waiting, between completion and reception, lose their phenomenological sharpness. The simulacrum does not erase these differences; it renders them increasingly difficult to perceive.

Understood in this way, the stakes are not merely philosophical. They are formative. A culture becomes what it practices. Capacities that are no longer exercised begin to atrophy. And what is replaced by a sufficiently convincing simulation may not disappear, but it can become inaccessible, not because it has been destroyed, but because it has ceased to appear as something that could be lived.

The bowl remains on the table. Nothing has replaced it. Still, it is becoming harder to see.

7. Radical Pause: From Critique to Practice

A diagnosis without direction risks becoming its own form of privation. If this argument were to end with the identification of the simulacrum, having named the stakes and traced their structural depth, it would remain, in an important sense, incomplete. Critique alone does not restore what has been obscured. Something like a counter-practice must at least become imaginable.

Any such response cannot be merely individual. Byung-Chul Han makes this point with characteristic force: the pathologies of contemporary achievement culture are not failures of personal discipline but consequences of systemic conditions that shape the field of possible action [22] (pp. 1–7). Practices of resistance that remain purely private, reducing screen time, cultivating mindfulness, withdrawing from technological systems, may have value, but they leave intact the structures that continuously reproduce the conditions

they seek to escape. What is required is not only withdrawal but the creation of alternative forms of life capable of sustaining different temporal and attentional habits.

Contemplative traditions offer precisely this kind of resource. Within the Christian monastic tradition, practices of silence, prayer, and liturgical time were developed in explicit recognition that the human person is not naturally disposed toward receptivity. Left to itself, attention disperses. It seeks stimulation, novelty, and control, what medieval writers described under the term *curiositas*, a restless movement of the mind that substitutes accumulation for understanding [23] (II-II, q. 167, a. 1). Against this tendency, contemplative practices cultivate a disciplined form of waiting: *lectio divina*, in which reading becomes a mode of listening; silent prayer, in which speech yields to attention; the Liturgy of the Hours, in which time itself is structured not by productivity but by recurring intervals of receptivity. These are not merely devotional exercises. They are practices of formation, sustained over time and within community, designed to make a different kind of attention possible. Teresa of Avila's account of the interior life is organized precisely around the progressive deepening of receptive capacity, the soul learning, through structured practice and communal support, to remain present to what it cannot control or complete [24] (Fourth Mansions).

Parallel resources emerge in secular contexts. Jenny Odell's account of "doing nothing" frames refusal not as disengagement but as a deliberate reorientation of attention away from the demands of productivity and toward forms of presence that cannot be easily instrumentalized [25] (pp. 55–60). Thich Nhat Hanh's articulation of mindfulness similarly emphasizes attention as a practice of sustained presence rather than efficient cognition, a way of inhabiting time that resists fragmentation [26] (pp. 64–70). Across these traditions, differences of metaphysics and vocabulary remain significant, but a shared insight emerges: receptive openness is not a default condition to which one returns once distractions are removed. It is a cultivated capacity, requiring repetition, structure, and, crucially, the support of others engaged in the same practice.

Institutional form therefore matters. If the surrounding environment is structured by acceleration, interruption, and the continuous solicitation of response, then the cultivation of receptivity requires spaces in which different norms are operative. Radical Pause: The UCCS Stillness Project proceeds from this premise. Its guiding intuition is simple but difficult to enact: that shared, purposive stillness can function as a counter-practice within a university setting otherwise organized around production, evaluation, and performance. The "Stillness Salon," as one expression of this project, gathers participants not to produce or exchange ideas but to inhabit a period of collective silence and to reflect, afterward, on what that silence discloses. The point is not withdrawal from thought but the suspension of its immediate instrumentalization.

The link between this practice and the technical analysis of Sections 3 through 5 is direct. What the Salon practices is exactly what the systems described there cannot represent or reward, the unresolved interval, the question allowed to remain a question, the response withheld without penalty. There is no loss function in the room. No token must come next. The silence is the substance of the gathering rather than latency to be managed, and the interval is entered for what it discloses rather than for what it prepares. In this sense the project is a deliberate inversion of the training environment this essay has described. Where the model is penalized at every step of its formation for failing to complete, participants practice, step by step, the capacity to leave things uncompleted, and discover that the capacity strengthens with use.

Such practices will appear, from within the dominant logic of optimization, as inefficient or even unintelligible. That appearance is not an objection but a confirmation of their necessity. A culture organized around completion will not spontaneously generate forms

of activity that resist completion. These must be intentionally created, sustained, and protected.

Protection raises a harder question, and the argument cannot end without facing it. Practices of pause are no longer foreign to the industry this essay has been describing. Screen-time dashboards now ship as features of the very devices they purport to discipline. Corporations sponsor mindfulness programs for their employees, meditation applications are venture funded and engagement optimized, and the silent retreat has become a product with a price point. If stillness can be absorbed this smoothly into the economy of attention, the objection runs, then radical pauses name a market segment rather than a resistance, and proposals like the one advanced here are already halfway to becoming amenities.

Han gives this objection its sharpest available form. The neoliberal regime, he argues, no longer confronts resistance from outside itself. It absorbs the practices of interiority, converting techniques of self-knowledge into techniques of self-optimization, so that the subject monitors and improves itself in the name of its own freedom, and mindfulness becomes an instrument for stilling the worker in order to render the worker more productive [27]. The objection is therefore serious, and it is answered by the essay's own central distinction. The corporate-sanctioned pause is the drained glass returning in the vocabulary of wellness. It is instrumentalized through and through, undertaken for the sake of restored productivity, measured by its contribution to performance, and discontinued the moment it stops paying. Its emptiness is privative, a deficit of output awaiting resolution by better output.

What distinguishes a genuine counter-practice is accordingly its structure of relation rather than its outward content, which may look identical in both cases, the same closed eyes in the same silent room. Three marks follow from the argument already made. A genuine counter-practice is non-instrumental; the stillness is entered for its own sake, and its value lies precisely in resolving toward nothing. It is communal and institutional in form, since, as Han himself insists, purely private discipline leaves systemic conditions intact [22] (pp. 1–7). And it is formative rather than technical; it aims at the slow reshaping of a capacity over time, on the model of the contemplative traditions surveyed above, rather than at a measurable improvement in any particular performance. Jenny Odell draws the same line from a secular direction when she distinguishes genuine refusal from the commodified self-care that the attention economy is happy to sell back to those it has exhausted [25]. By these marks, a screen-time feature and a Stillness Salon can be told apart even on the days when they look alike, and Radical Pause is answerable to the marks as fully as any corporate program is. The project earns the name counter-practice only for as long as its stillness resolves toward nothing.

What is at stake in such efforts is modest in scale but significant in implication. The aim is not to reject technological systems outright, nor to recover a pre-digital form of life that is no longer available. It is, rather, to preserve and cultivate a human capacity that those systems cannot provide, the capacity to remain with what does not resolve, to attend without immediately completing, to receive without converting what is received into output.

Such a capacity cannot be engineered. It can only be practiced.

The bowl remains where it has been throughout. It does not ask to be used. It does not solicit attention. It does not optimize its presence. It simply offers a form: open, patient, available. Beneath the icon of Our Lady of Perpetual Help, whose own posture is one of total receptive availability, *fiat mihi secundum verbum tuum* (let it be done to me according to your word), the bowl has simply been what it is.

Learning to see it again may require more than noticing. It may require learning, together, how to stop filling.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data sharing is not applicable to this article. No new data were created or analyzed in this study.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Marion, J.-L. *Being Given: Toward a Phenomenology of Givenness*; Kosky, J.L., Trans.; Stanford University Press: Stanford, CA, USA, 2002.
2. Derrida, J. *Given Time: I. Counterfeit Money*; Kamuf, P., Trans.; University of Chicago Press: Chicago, IL, USA, 1992.
3. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention Is All You Need. In *Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Lane Red Hook, NY, USA, 2017; Volume 30.
4. Weil, S. *Waiting for God*; Craufurd, E., Trans.; Harper Perennial: New York, NY, USA, 2009.
5. Weil, S. *Gravity and Grace*; Crawford, E., von der Ruhr, M., Trans.; Routledge: London, UK, 2002.
6. Tay, Y.; Dehghani, M.; Bahri, D.; Metzler, D. Efficient Transformers: A Survey. *ACM Comput. Surv.* **2022**, *55*, 1–28.
7. Zhao, W.X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.A.; et al. Survey of Large Language Models. *arXiv* **2023**, arXiv:2303.18223.
8. Bender, E.M.; Gebru, T.; McMillan-Major, A.; Shmitchell, S. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT' 21), Virtual Event, Canada, 3–10 March 2021.
9. Han, B.-C. *In the Swarm: Digital Prospects*; Butler, E., Trans.; MIT Press: Cambridge, MA, USA, 2017.
10. Bergson, H. *Time and Free Will: An Essay on the Immediate Data of Consciousness*; Pogson, F.L., Trans.; Dover: Mineola, NY, USA, 2001.
11. Han, B.-C. *The Scent of Time: A Philosophical Essay on the Art of Lingering*; Steuer, D., Trans.; Polity Press: Cambridge, UK, 2017.
12. Pope, R.; Narang, S.; Klein, D.; Levine, S. Efficiently Scaling Transformer Inference. *arXiv* **2022**, arXiv:2211.05102.
13. Wan, Z.; Wang, X.; Liu, C.; Alam, S.; Zheng, Y.; Liu, J.; Qu, Z.; Yan, S.; Zhu, Y.; Zhang, Q.; et al. Efficient Large Language Models: A Survey. *arXiv* **2023**, arXiv:2312.03863.
14. Inan, M.; Lehman, J.F.; Black, A.W. Better Slow than Sorry: Introducing Positive Friction for Reliable Dialogue Systems. *arXiv* **2025**, arXiv:2501.17348.
15. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models Are Few-Shot Learners. In *Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Lane Red Hook, NY, USA, 2020; Volume 33.
16. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E. On the Opportunities and Risks of Foundation Models. *arXiv* **2021**, arXiv:2108.07258.
17. Hoffmann, J.; Borgeaud, S.; Mensch, A.; Buchatskaya, E.; Cai, T.; Rutherford, E.; Casas, D.D.L.; Hendricks, L.A.; Welbl, J.; Clark, A.; et al. Training Compute-Optimal Large Language Models. *arXiv* **2022**, arXiv:2203.15556.
18. Amershi, S.; Weld, D.; Vorvoreanu, M.; Fournery, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P.N.; Inkpen, K.; et al. Guidelines for Human-AI Interaction. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI' 19), Glasgow, UK, 4–9 May 2019.
19. Liao, Q.V.; Vaughan, J.W. AI Transparency in Practice: Design and Evaluation of Human-AI Interaction. *Commun. ACM* **2023**, *66*, 68–77.
20. Dao, T.; Fu, D.Y.; Ermon, S.; Rudra, A.; Re, C. FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness. In *Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Lane Red Hook, NY, USA, 2022; Volume 35.
21. Baudrillard, J. *Simulacra and Simulation*; Glaser, S.F., Trans.; University of Michigan Press: Ann Arbor, MI, USA, 1994.
22. Han, B.-C. *The Burnout Society*; Butler, E., Trans.; Stanford University Press: Stanford, CA, USA, 2015.
23. Aquinas, T. *Summa Theologica*; Fathers of the English Dominican Province, Trans.; Benziger Bros.: New York, NY, USA, 1947.
24. Teresa of Avila. *The Interior Castle*; Peers, E.A., Trans.; Image Books: New York, NY, USA, 1961.
25. Odell, J. *How to Do Nothing: Resisting the Attention Economy*; Melville House: Brooklyn, NY, USA, 2019.

26. Nhat Hanh, T. *The Miracle of Mindfulness*; Beacon Press: Boston, MA, USA, 1999.
27. Han, B.-C. *Psychopolitics: Neoliberalism and New Technologies of Power*; Butler, E., Trans.; Verso: London, UK, 2017.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.