BREAKING THE IMITATION BARRIER

Why Today's Artificial Intelligence Imitates, and What It Will Take for Machines to Create

A White Paper for Business Leaders

Rowland Chen, Roger B. Dannenberg, Bhiksha Raj, and Rita Singh

School of Computer Science and Language Technologies Institute, Carnegie Mellon University

Adapted from a peer-reviewed paper presented at the 11th International Conference on Computational Creativity (Association for Computational Creativity, 2020).

The Idea in Brief

Generative artificial intelligence has dazzled the business world by painting in the style of the masters, impersonating voices, and drafting prose on demand. We argue these systems share a hidden ceiling. They imitate; they do not create. We call this ceiling the imitation barrier, and we believe it defines the boundary of competitive advantage available from current AI.

This paper introduces the Insight-Knowledge Object model (IKO, a hierarchical architecture of eleven levels of knowledge and ten insight processes that act upon them) as a framework for understanding why the barrier exists and how it might be broken. The model holds a practical warning for executives: systems that skip levels of knowledge, as today's deep learning systems do, collapse short of genuine creation. We describe a working test of the framework, a machine designed to improvise music in real time alongside human performers, and we draw out the implications for any firm betting its innovation agenda on AI.

Imitation Is Not Innovation

Since Joseph Jacquard punched wooden cards to weave fabric designs in 1801, computing machinery has excelled at one thing above all: imitation. Recent advances in machine learning extend that tradition with remarkable polish. Style transfer systems approximate the brushwork of select painters. Neural networks impersonate voices with unsettling fidelity. Large models generate text that reads as fluent and confident.

Yet creativity has a stricter definition. Scholars from Margaret Boden onward define a creative idea as one that is novel, surprising, and valuable. Nothing truly novel and no surprises result from imitation. Alan Turing asked in 1950, “Can machines think?” We ask the question that now matters for science and for commerce: can machines create?

The distinction is not academic. A firm that deploys generative AI is purchasing imitation at scale. Imitation at scale is valuable; it compresses cost and accelerates production. But it does not originate. The breakthrough product, the unexpected strategy, the category-defining campaign: these remain the province of human inspiration. Understanding precisely where the machine stops and the human begins is, we contend, a first-order question for any leadership team allocating capital to AI.

A Hierarchy of Knowledge: The Insight-Knowledge Object Model

Not all knowledge is created equal. The Insight-Knowledge Object model proposes that human cognition ladders upward through eleven levels of knowledge objects, each generated by an insight process acting on the level beneath it (Figure 1). The hierarchy begins in a state we call void, in

which even the acknowledgement of nothing does not exist, and culminates in creation, generated by inspirational insight.

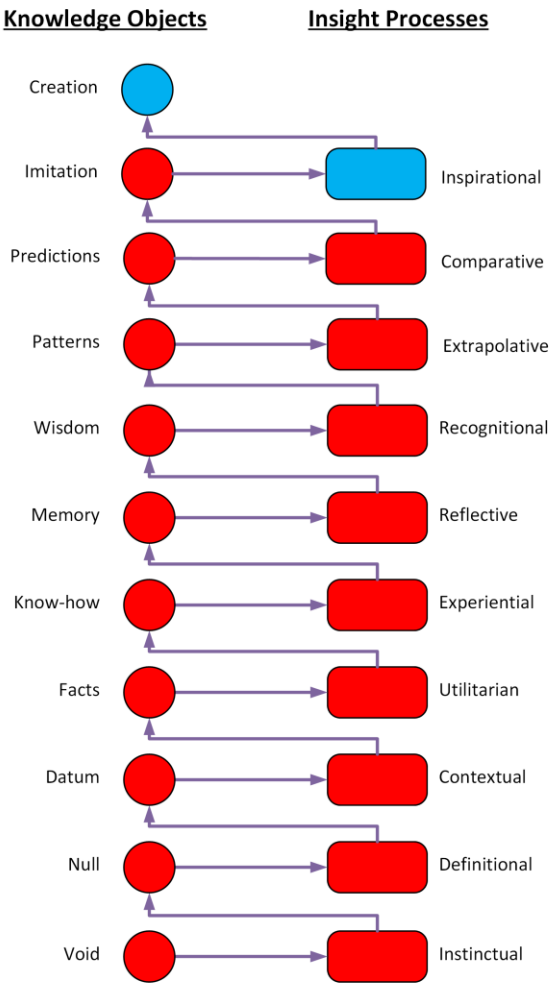

Figure 1: The Insight-Knowledge Object model, a hierarchical architecture for human cognition. Knowledge objects appear on the left; the insight processes that generate them appear on the right.

The ascent proceeds as follows. Instinctual insight acts on the void to produce null, a primal awareness of one's environment. Definitional insight labels the unnamed, producing data. Contextual insight gives data meaning, producing facts. Consider the word coffee. As a datum it carries no reference. In the commodities market it becomes a traded good; in food service it becomes a beverage. Utilitarian insight then produces know-how (knowledge of use): how to trade coffee futures, or how to prepare an espresso. Experiential insight converts executed know-how into memories, the way a barista remembers how much foam belongs atop a latte.

The upper levels grow more abstract. Reflective insight acts on sets of memories to produce wisdom, the understanding of why customers order lattes rather than merely how to make them.

Recognitional insight connects pieces of wisdom into patterns, the domain of data science. A pattern can join subjects as distant as a barista perfecting a pour and the swordsman in Kurosawa's Seven Samurai who dedicates his life to mastery; the connection is the insight. Extrapolative insight converts patterns into predictions, the domain of statistics and forecasting. Comparative insight tests predictions against references to produce imitations, the domain of today's machine learning and of generative adversarial networks (GANs, paired neural networks that refine output by competition).

And there the modern machine stops. The final step, inspirational insight acting on imitation to produce creation, has not been designed. A master barista who conceives an original design in steamed milk is not mimicking prior art (Figure 2). She is creating. No deployed system does the same.

Figure 2: An inspired latte design. The master barista creates rather than copies. (Photo credit: 123rf.com)

Readers familiar with knowledge management will recognize a forerunner in the DIKW pyramid (data, information, knowledge, wisdom), a model that has not evolved since its introduction late last century. The IKO model is intended primarily for artificial creative intelligence, yet it offers a collateral benefit: a richer methodology for organizing knowledge inside the firm, one that distinguishes know-how from facts and wisdom from patterns. We suggest this is one of several *better practices* (we prefer the term to "best practices," which implies a finality that practice rarely earns) available to leaders who take knowledge architecture seriously.

The Imitation Barrier

The model can be written compactly. Each knowledge object is a function of the object one level beneath it, and the full system is a composition of all ten insight functions. The mathematics matter

less than the consequence: if any function in the chain is skipped, the system collapses, and the output falls short of creation.

Here lies the diagnosis. Today's most celebrated AI systems act directly on data, level three of the hierarchy, and leap to imitation, level ten. Facts, know-how, memories, and wisdom are bypassed entirely. We hypothesize that this level skipping is precisely why current approaches cannot break through to creation (Figure 3). The boundary of the state of the art sits at imitation, and brute computational force does not move it.

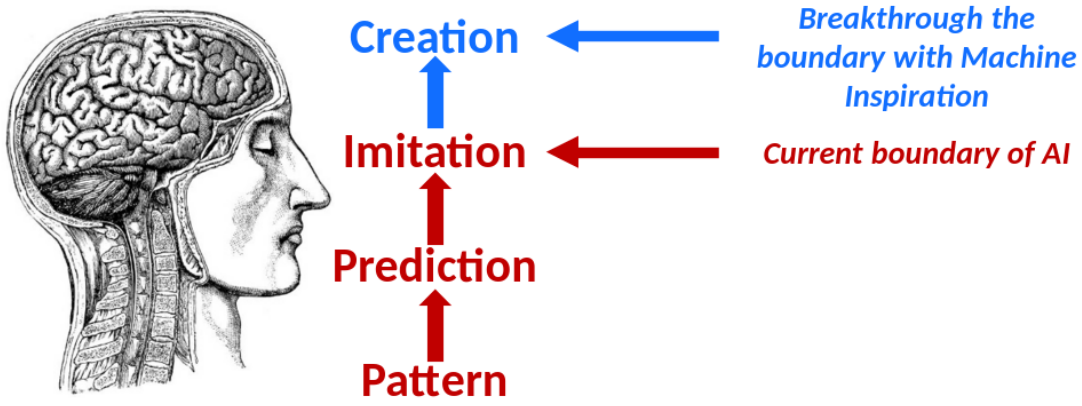

Figure 3: The imitation barrier. Current AI ascends from pattern to prediction to imitation and stops. Breaking through to creation requires machine inspiration.

Practitioners at the frontier concede the point. Jerome Pesenti, then Vice President of Artificial Intelligence at Facebook, observed in 2019 that deep learning, honestly assessed, has significant limitations. John Smith, an IBM Fellow at IBM Research, reached a similar conclusion from IBM's experiments with machine-generated movie trailers: deep learning is not the answer to creativity.

When a human creates, inspiration does not happen in a vacuum. All knowledge in the brain comes into play. A system that omits whole strata of knowledge is structurally incomplete, and structures with missing levels share a familiar fate (Figure 4).

Figure 4: Skip a level and the system collapses like a house of cards. (Photo credit: 123rf.com)

The Framework in Practice: JerryBot

A framework earns its keep in application. Our test case at Carnegie Mellon University is music improvisation, an unambiguous instance of human creativity: novel, surprising, and valuable enough that improvising artists build devoted audiences worldwide. When a beloved artist dies, the new improvisations and the fan experiences that accompany them disappear forever. Our goal is an Artificial Creative Intelligence machine that performs alongside human musicians in real time, producing music indistinguishable from what the artist might have played. One can call this the creation game, with apologies to Turing.

The system, called JerryBot, focuses on the late guitarist Jerry Garcia of the Grateful Dead. Its design integrates four insight functions from the IKO model, recognitional, extrapolative, comparative, and inspirational, into a single architecture (Figure 5). Living musicians supply the audio input; the machine listens, predicts, compares, and, in the component still under fundamental research, draws inspiration to improvise a lead guitar line of its own.

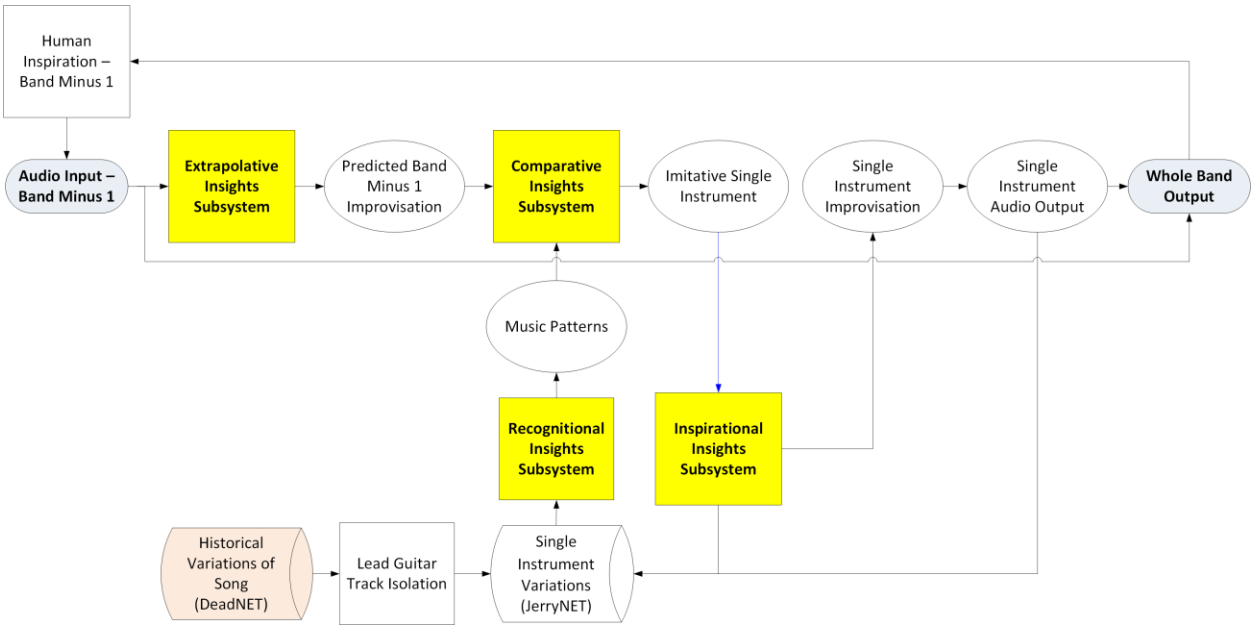

Figure 5: The Artificial Creative Intelligence machine: JerryBot design (Chen, 2020). The four highlighted subsystems implement recognitional, extrapolative, comparative, and inspirational insight functions.

Crucially, the design skips no levels. Its foundation is DeadNET, a database extracted from the Internet Archive containing recordings of nearly 2,500 Grateful Dead concerts performed between 1965 and 1995, approximately 7,500 hours of music (Figure 6). Within DeadNET sit more than six hundred variations of the song “Playing in the Band” alone. Embedded in three decades of performance are the band's musical data, facts, know-how, memories, and wisdom, the very levels current systems discard.

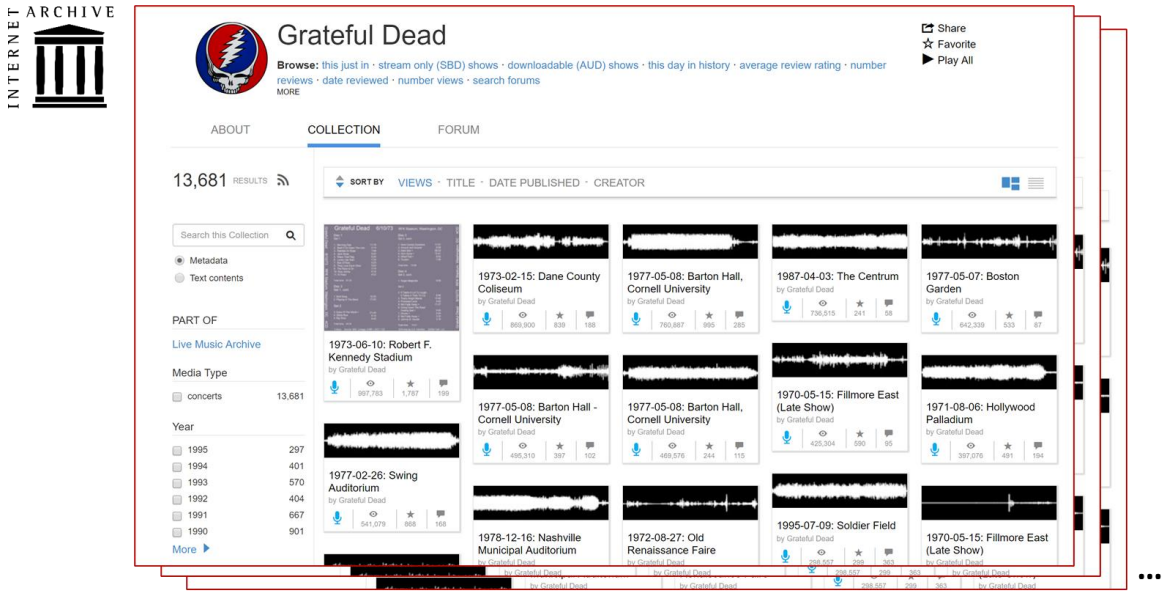

Figure 6: DeadNET, extracted from the Grateful Dead collection of concert recordings on the Internet Archive.

We have adopted a music-as-language paradigm to leverage tools developed in Carnegie Mellon's Language Technologies Institute, including a bag-of-words model and vector quantization of audio files (Figure 7). Continuing work addresses the isolation of Garcia's guitar track from full-band recordings, the modeling of musical conversation within an ensemble, and the hardest problem of all: how a machine can assess the pleasantness, the value, of what it has created.

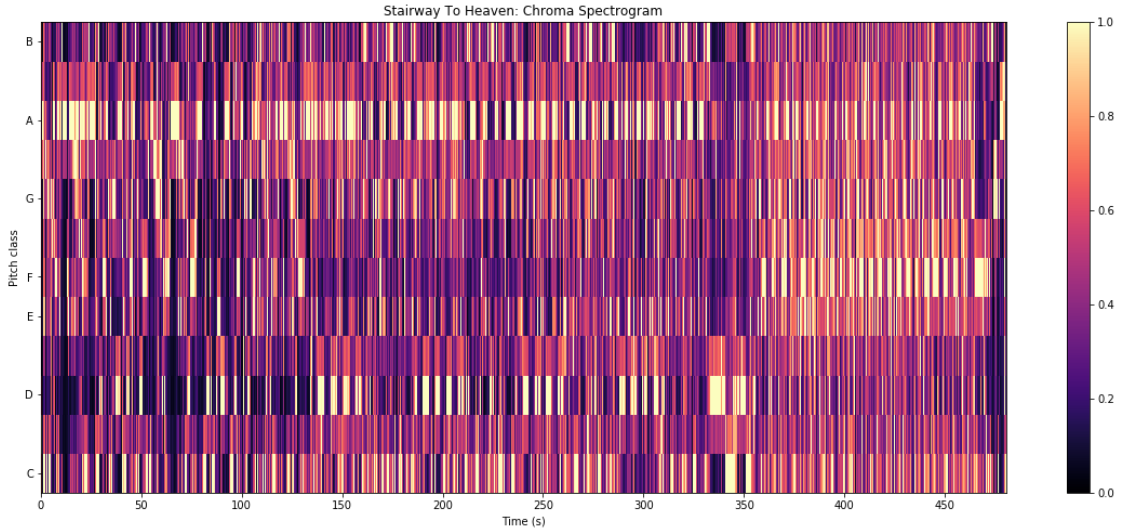

Figure 7: Log mel spectrogram from work in process (Raj, Agarwal, and Raj, 2019).

Beyond Music: Where Creative Machines Will Matter

Music is the proving ground, not the destination. Consider two further cases. The first is the haiku problem. A haiku's first two lines set a tone; the third must deliver a surprising, emotionally satisfying close that diverts the reader from any easy extrapolation. Machines can count syllables. The elusive novel close, the moment of inspiration, remains beyond them.

The second case carries commercial and ethical weight: the no-win self-driving car problem. Autonomous vehicles integrate robotics, image recognition, and deep learning, but their responses are limited to learned scenarios. Confront a vehicle with an unanticipated, unlearned, no-win situation and the limits of imitation become matters of life and death. We contend that currently only a human can make an acceptable conscious choice in such a scenario, and that acceptability itself differs across cultures. An Artificial Creative Intelligence, developed first for something as benign as music, could one day supply the real-time creative judgment required to minimize physical, emotional, ethical, and societal damage.

What This Means for Leaders

Three implications follow for the executive agenda. First, calibrate expectations. Generative AI deployed today buys imitation, and imitation is genuinely valuable for cost, speed, and scale. But a firm should not mistake fluent mimicry for an innovation engine. Strategy work that assumes machines will originate the next breakthrough is, for now, built on a missing level.

Second, audit the knowledge architecture of the firm itself. The IKO model describes human cognition, yet it doubles as a diagnostic for organizations. Many firms hoard data and patterns while letting know-how walk out the door, and wisdom goes uncaptured. A knowledge strategy that spans every level of the hierarchy positions a firm to exploit creative machines if and when the barrier breaks.

Third, watch the interdisciplinary frontier. Breaking the imitation barrier will require research across mathematics, computer science, neuroscience, physiology, psychology, and perhaps philosophy. Open questions abound. Are matrix algebra, calculus, statistics, and probability sufficient to push the boundary into creative insight processing, or is new mathematics needed? Will quantum computing or analog networks of solid-state neurons play a role? The firms that monitor these questions will not be surprised by the answers.

Questions for the Reader

We close, as we believe every serious analysis should, by turning the inquiry toward you. How much of your AI investment is buying imitation, and how much is premised on creation that no machine yet delivers? Which levels of the knowledge hierarchy does your firm capture systematically, and which evaporate with every retirement and resignation? When the imitation barrier finally breaks, and we believe it will, what will your firm do in the first hundred days? And the question beneath them all, the one Turing would recognize: can machines create, and is your firm ready for either answer?

About this paper

This white paper is adapted from Rowland Chen, Roger B. Dannenberg, Bhiksha Raj, and Rita Singh, "Artificial Creative Intelligence: Breaking the Imitation Barrier," in Proceedings of the 11th International Conference on Computational Creativity, Association for Computational Creativity, 2020. Audio Intelligenzia provided in-kind support for the original work. Full citations to the scholarly literature referenced here, including Boden (2004), Turing (1950), Rowley (2006), Nyholm and Smids (2016), and others, appear in the original paper.

Artificial Creative Intelligence

.