

Artificial Creative Intelligence: Breaking the Imitation Barrier

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

School of Computer Science and Language Technologies Institute

Carnegie Mellon University

Pittsburgh, Pennsylvania 15213 USA

rowlandc@andrew.cmu.edu, rbd@cs.cmu.edu, bhiksha@cs.cmu.edu, rsingh@cs.cmu.edu

Abstract

Not all knowledge is created equal. A hierarchical architecture is a method to classify knowledge for use in the field of human cognition and computational creativity. This paper introduces an Insight-Knowledge Object (IKO) model as a framework for Artificial Creative Intelligence (ACI), a step forward in the pursuit of replicating general human intelligence with computing machinery. To achieve ACI, it is hypothesized that a fundamental rethinking of the architecture of human cognition and knowledge processing is required. One possible novel architecture could be the IKO model. The authors include a description of on-going work at Carnegie Mellon University that applies the IKO model in practice with an artificial music improvisation embodiment.

Introduction

Since the dawn of computing machinery (Jacquard’s first use of wooden punch cards for looms to imitate and weave fabric designs in 1801, from Zimmerman, 2017), the challenge of achieving true machine creativity remains elusive (IBM, 2019; Olenik, 2019). Creativity has been defined as an idea that is novel, surprising, and valuable (Boden, 2004; Abraham, 2018; Runcon, 2012; Stein, 1953). Recent advances in artificial intelligence, machine learning in particular, approach the successful imitation of the styles of select painters through style transfer in the visual arts (Zhou, Z., *et al.*, 2019). Advances have also been achieved in voice impersonation (Andabi, 2017; Gao, 2018). However, the authors assert true creations do not emerge from style transfer. Nothing truly novel and no surprises result from imitation. Similar to Turing’s question from 1949, “Can machines think?” (Turing, 1950) the authors approach their ultimate

question, “Can machines create?”

A fundamental re-thinking of human and artificial creative processes and corresponding mathematics and computing machinery are indicated to break through an imitation barrier. This paper introduces an Insight-Knowledge Object (IKO) model as a framework for human creativity and Artificial Creative Intelligence (ACI), a step forward in the pursuit of replicating general human intelligence with computing machinery. The IKO model represents a novel approach to computational creativity and computing machine design.

Of importance to consider are the ethical and moral implications of developing a machine that can create new thoughts, ideas, plans of action, conversations, science works, and works of art as would a human. While one clear answer does not exist, embedded in the question facing scientists and engineers, “*Should* machines create?” are issues and concerns of ethics and morality.

The authors hope the Insight-Knowledge Object model proves useful to those advancing the science of human intelligence and the research driving towards computational creativity.

Insight-Knowledge Object Model

A proposed Artificial Creative Intelligence framework, the Insight-Knowledge Object (IKO) model (Figure 1), builds on Chen’s prior work (Chen, 2009). Proposed here is a way of modeling the human thought process that can be used to shape the development of ACI machines. It is important that each level in the IKO hierarchy be present in the process of human thinking as well as in artificial attempts at reproducing that process. Knowledge objects are on the left-hand side, and insight processes appear on the right. In the IKO model, as insight processes and knowledge objects ladder up

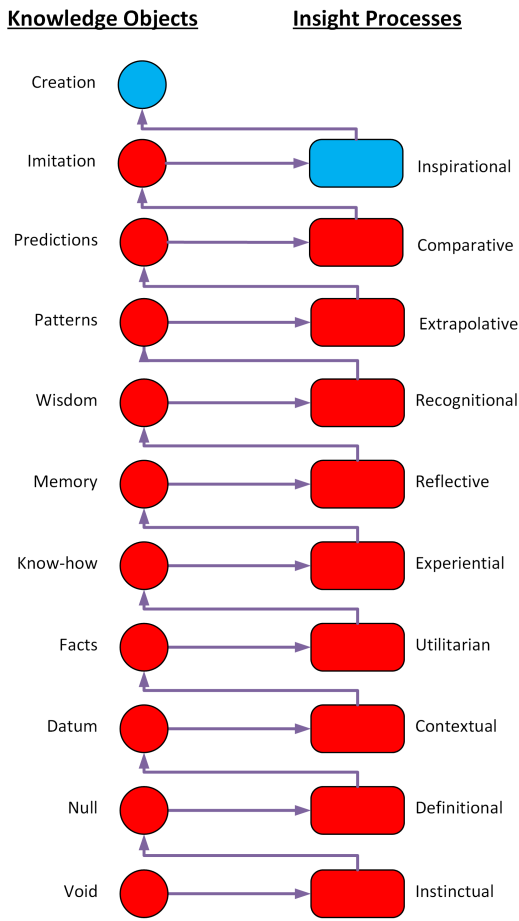


Figure 1: Insight-Knowledge Object model representing a hierarchical architecture for human cognition.

through the hierarchy, increasing levels of cognitive sophistication are reached. The IKO model has eleven levels of knowledge objects and ten levels of insight processes acting upon the knowledge objects. The IKO model begins with a state of knowledge object known as void. In this state, even the acknowledgement of nothing does not exist. At the very top of the IKO model is the knowledge object creation generated by an inspirational insight process.

Descriptions of Each Insight Process

Instinctual insight acts on the void to generate null, the first true layer of knowledge. Null emerges from void as primal instincts create an awareness of one’s environment. This level of knowledge is deemed null because at this level, there is at least consciousness of existence or non-existence. In the void, even consciousness does not exist.

Definitional insight acts on the null to elevate knowledge into data. At this level, the insight process gives definition to objects and actions in the null. Definitional insight labels a collection of unnamed and unidentified things so that

distinctions are drawn between them. Each object is now defined and becomes a datum.

Contextual insight acts on data to generate facts in the next layer of the hierarchy. Facts represent a richer and fuller set of knowledge than pure data. For example, if one takes the word coffee as a datum there is no context for reference. Given some context such as the commodities trading market, coffee takes on the meaning of a traded good. If food service is the context, then coffee takes on the meaning of a beverage. Contextual insight allows distinctions to be made between data to create different facts.

Utilitarian insight acts on facts to generate know-how, how an object is used and for what purpose. In our coffee example, utilitarian insight emerges to provide the know-how for what to do with coffee. In the commodities market context, know-how would be how to trade coffee on the spot or futures markets. In the food service context, know-how would be how to prepare coffee for consumption. Without utilitarian insight, coffee has no real value. Simply speaking, utilitarian insight provides knowledge of use.

Experiential insight acts on know-how to generate memories. The execution of know-how generates experiences that can be remembered and used in the future. Following the coffee example, experience in making coffee enables a barista to remember how much foam to put on top of a latte.

Reflective insight acts on memories to generate wisdom. Reflection works on a meta-plane of thinking and takes on a new layer of abstraction in the knowledge hierarchy. Insights are not simply generated on single points of execution but a set of memories. For example, remembering how to make a latte is a memory but digging deep to understand why people order lattes takes reflection. Wisdom emerges as a person can take a step back to reflect and learn from prior thoughts, decisions, and actions.

Recognitional insights act on wisdom to generate patterns. This insight function is in the realm of data science and data analytics. Recognizing patterns links related or unrelated pieces of wisdom to generate knowledge that would not emerge otherwise. For example, connecting the preparation of a perfect latte to the film “Seven Samurai” (Kurosawa, 1954) in which one of the samurai has dedicated his whole life to perfect his skills as a swordsman represents connecting two topics that on the surface are completely unrelated. A pattern emerges with the recognition that these humans continually strive for mastery in their respective fields. Recognitional insights create connections that drive thinking further and produce patterns.

Extrapolative insights act on patterns to generate predictions. This insight function is in the realm of statistics and probability. Making predictions links related or unrelated patterns to generate knowledge that would not emerge otherwise. Extrapolative insights produce predictions used for weather forecasting, social media user preference, and the serving up of relevant advertising. Again, in our coffee example, extrapolative insight is required to predict the demand for lattes in a coffee shop during the course of a day.

Comparative insights act on predictions to generate imitations. This insight function is in the realm of machine learning and implementations such as generative adversarial networks (GANs). Predictions are compared against a reference with the goal of achieving an imitation that most closely matches the reference. Much work and many examples of painter style matching, voice impersonation, and literature already exist. This level of the knowledge hierarchy is the present boundary of today’s state-of-the-art deep learning techniques. One could surmise researchers and scientists have hit an imitation barrier. For lattes, skilled baristas can copy fanciful milk designs on the surface of lattes based on prior examples of a master barista’s work. This is latte style transfer.

Inspirational insights act on imitations to generate creations. This insight function has not yet been designed and requires fundamental research across the disciplines of neuroscience, physiology, psychology, and computer science. The proposed ACI framework surpasses the imitation level by using unique, to-be-developed computing machines and computational algorithms, simulating human inspirational insights, to output creations. A master barista thinks up and executes unique and novel designs for lattes (Figure 2). She is not simply mimicking prior art. The inspiration process of the proposed IKO model breaks through the imitation barrier. At this level of the IKO hierarchy, the challenge of Artificial Creative Intelligence could be solved.



Figure 2: Inspired latte design (Photo credit: 123rf.com)

Related Work – the DIKW Pyramid

In information management and business technology, knowledge management was a new approach for data storage and data processing that was introduced late last century. A hierarchical model was developed at the time to capture, store and retrieve documents (origin unknown, Sharma, 2008). Data, information, knowledge, and wisdom comprise the model, also known as the DIKW pyramid (Zeleny, 2005; Rowley, 2006). In the knowledge management field of information technology, the DIKW pyramid has not evolved since its introduction (Elliott, 1934). A collateral benefit of the new IKO model, primarily intended for artificial creative intelligence applications, is a potential new methodology for knowledge management in business information management and information technology.

The IKO Model as Functions

The increasing levels of sophistication in the IKO model can be written as a set of functions in which insight processes act upon knowledge objects (variables) of the form:

Knowledge object = Insight function acting on the next lower order of knowledge object

$$KO_i = f(KO_{i-1})$$

The increasing levels of sophistication in the IKO model are symbolically presented below where a function is an insight process. The exact nature of the functions has not yet been discovered. The use of the “ f_i (KO)” notation is intended to drive the thinking of those seeking ways to represent creativity and perhaps, to provide a way to develop the mathematics of creativity.

void	=	initial condition
null	=	f_1 (void)
datum	=	f_2 (null)
fact	=	f_3 (datum)
know-how	=	f_4 (fact)
memory	=	f_5 (know-how)
wisdom	=	f_6 (memory)
pattern	=	f_7 (wisdom)
prediction	=	f_8 (pattern)
imitation	=	f_9 (prediction)
creation	=	f_{10} (imitation)

Therefore, creation is a function of imitation, prediction, pattern, wisdom, memory, know-how, fact, datum, null, and void.

The authors postulate the system function, F , relies on nine nested, or composite, functions. If any function is skipped, the system collapses, and output falls short of a creation. F can thus be written as:

$$F = f_{10} \circ f_9 \circ f_8 \circ f_7 \circ f_6 \circ f_5 \circ f_4 \circ f_3 \circ f_2 \circ f_1$$

Breaking the Imitation Barrier

“Deep learning and current AI, if you are really honest, has a lot of limitations.”

Jerome Pesenti, Vice President of Artificial Intelligence, Facebook (December 4, 2019)

“Creativity and where the authors started exploring with the (Disney movie trailer) is fascinating because deep learning isn’t the answer to creativity.”

John Smith, IBM Fellow at IBM Research (2019)

Current attempts at computational imitation rely on machine learning techniques that use algorithms acting on training and testing data in a resource-intensive approach. In

the context of the proposed IKO model, generative adversarial networks (GANs) produce imitations (level ten) acting directly on data (level three), thus skipping several levels. But imitations are not creations. The authors hypothesize the existence of an imitation barrier that limits current approaches to level ten of the model (Figure 3). The authors also hypothesize that as a result of level skipping in the context of the IKO framework, attempts at artificial creativity will not be able to break through an imitation barrier to reach creations.

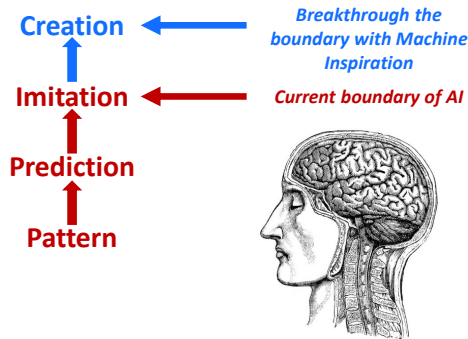


Figure 3: The Imitation Barrier

Computational creativity occurs in natural brain function, and the proposed IKO hierarchy suggests a model for human thought processes. When a human creates, inspiration does not happen in a vacuum. All knowledge in the brain comes to play. For example, facts, know-how, memories, and wisdom are needed but are left out of current imitative computing. Refer to the system function, F , that is a composite of all insight functions. If any function is skipped, the system collapses.



Figure 4: House of cards collapsing (Photo credit: 123rf.com)

In order to validate or refute the level-skipping failure hypothesis, additional research and exploration are required, using an interdisciplinary approach across the fields of mathematics, computer science, computer engineering,

neuroscience, physiology, psychology, and perhaps, philosophy.

JerryBot: the IKO Model in Practice

In various music genres, creative improvisation on a single or multiple instruments delivers music that is pleasing to the listener. The music genres include rock, jazz, jazz fusion, classical, blues, rhythm & blues, bluegrass, Indian, and others. The improvisation technique of one member, or more members, of a musical group creates new, real-time compositions during a live performance or in a music studio. Music improvisation is an example of human creativity – novel, surprising, and valuable (Boden, 2004; Abraham, 2018; Runcon, 2012; Stein, 1953). The sounds created are so valuable (pleasing) that a music artist develops a fan base comprised of music lovers around the world. The improvised music is created through inspirational insights based on the intelligence, experience, and technique of the music artist or artists.

When a highly skilled and highly loved music artist passes away, his or her new improvisational music creations and corresponding fan experiences disappear forever. An Artificial Creative Intelligence machine fills the emptiness left behind by a dead artist by recreating the improvisation style and reproducing the instrument sounds of the specific artist. The music fan’s pleasurable listening experience with the late artist and his or her band is extended. The goal of the machine is to perform alongside human musicians in real-time, and this hybrid human-artificial band’s music is indistinguishable from any new music that could have been performed if the artist were not dead. One can call this the creation game, with apologies to Turing (Turing, 1950). A version of an artificial music improvisation machine design appears in Figure 5.

In this system, called JerryBot (Chen, 2020), multiple living musicians improvise with the ACI machine, similar to a completely human band. On-going work at Carnegie Mellon University focuses on a specific musician, the late guitarist Jerry Garcia of the rock group The Grateful Dead. A database containing recordings of all live concert performances of the Grateful Dead has been captured from the Internet Archive (Internet Archive, 2019). These data, which the authors call DeadNET (Figure 6), represent nearly 2,500 concerts, performed between 1965 and 1995, and approximately 7,500 hours of music. The data comprise the Historical Variations of Song. For example, within DeadNET, there are over six hundred variations of the song “Playing in the Band.”

In the block diagram, one can see recognitional, extrapolative, comparative, and inspirational insight functions highlighted in yellow. These four insight functions are integrated into a single system in an attempt at developing an ACI machine. Additionally, embedded in the recordings captured in DeadNET are the band’s musical data, facts, know-how, memories, and wisdom accumulated over 30 years of per-

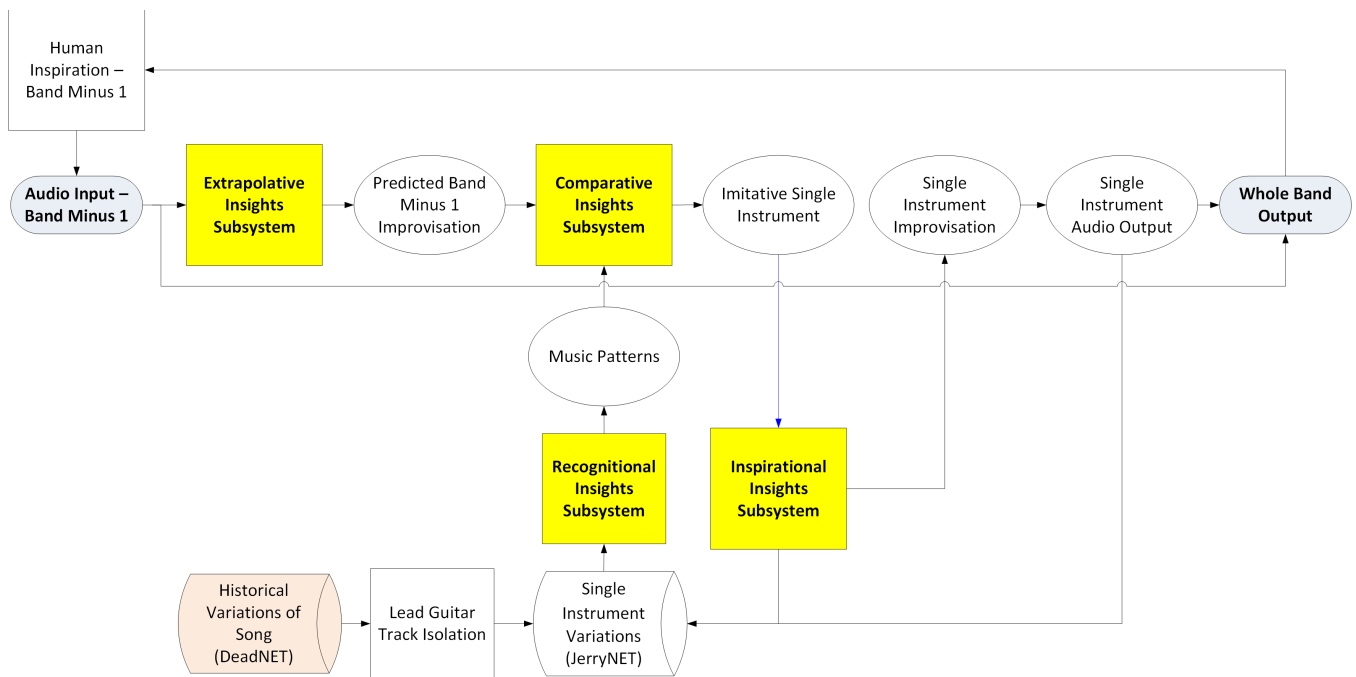


Figure 5: Artificial Creative Intelligence machine: JerryBot design (Chen, 2020)

formance. With DeadNET and the insight functions of recognition, extrapolation, comparison, and inspiration, the ACI machine under development includes all levels of knowledge object and insight function. No levels are skipped. The machine holistically brings together a whole-brain approach to computational creativity.

entire band to create JerryNET perhaps modeled after the neural basis of auditory attention used for speaker isolation (Geravanchizadeh, 2020), discovering the neural basis of inspirational insight, modeling the music conversations that occur when a jam band assembles and plays, and formulating how a machine can self-assess the value of computationally-created music (pleasantness of the output from a hybrid human-machine band). Much work remains.

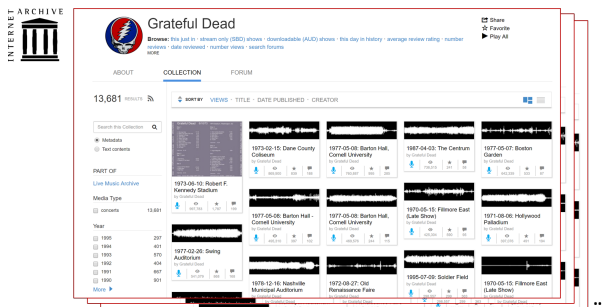


Figure 6: DeadNET extracted from the Grateful Dead collection of concert recordings on the Internet Archive.

The authors have adopted music as a language paradigm to leverage existing solutions and tools developed in the Language Technologies Institute at Carnegie Mellon University as well as open-source software. In process are the use of a “bag of words” model and vector quantization of .wav files (Figure 7, Raj, 2019). Continuing work is intended to address the other components of the artificial music improvisation machine, for example, the isolation of Jerry Garcia’s lead guitar track from the recordings of the

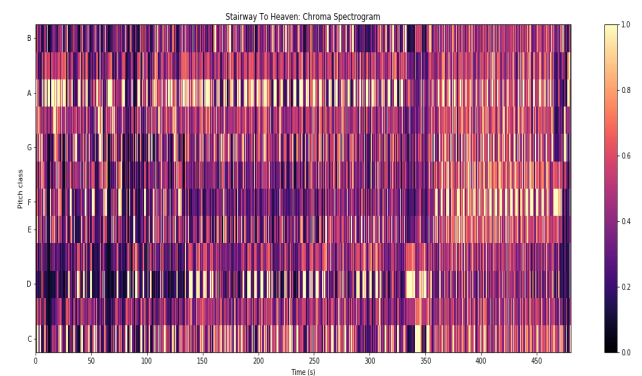


Figure 7: Log mel spectrogram from work-in-process (Raj, B., Agarwal, Raj, T., 2019)

Other Use Cases

Beyond music, other potential applications for Artificial Creative Intelligence include use cases in language arts (the *haiku* problem) and autonomous vehicles (control and decision-making in unlearned situations).

The *Haiku* Problem

5-7-5. That is the syllable count for the Japanese three-line poem style known as *haiku*. In *haiku*, the first two lines set the tone of the poem and the emotions of the reader. The third line is a surprising closure. For example, below is a *haiku* by Edo-period (18th century A.D.) Japanese artist Katsushika Hokusai:

*I write, erase, write,
Erase, re-write again, then
A red poppy blooms*

The final line of the *haiku* is a surprise and an emotionally satisfying ending created by Hokusai in a moment of inspiration. The third line also diverts the reader from extrapolation and a prediction of what might conclude the poem.

The *haiku* problem is for a machine to generate a third line that is both surprising and pleasing to the reader. Attempts at machine-written *haiku* have been able to achieve the correct syllable count (for example, in Aguiar, 2019). However, further work is required to achieve an elusive novel close.

The No-Win Self-Driving Car Problem

A continuous flow of conscious choices exists in driving situations. Even in relatively stable conditions, for example, cruising down the highway, humans choose speed, acceleration/deceleration, and direction. Add in entertainment, and the choices grow. Currently, autonomous vehicles integrate robotics, image recognition, comparative insight, and deep learning. However, a vehicle's response to turbulence introduced into the environment is limited to learned scenarios. But what about a situation unanticipated and unlearned by the autonomous systems?

An example of this would be a scenario in which multiple externalities are introduced into the moving vehicle environment to create a no-win situation. How would a self-driving car, with a passenger, make a decision between crashing into an 80-year old man, a pregnant woman, or a toddler or driving off a cliff (there being no other alternatives)? The controller of the vehicle has several choices to make, be the controller a human or an AI system. As Nyholm and Smids (Nyholm, 2016) have written, this problem is being addressed in the development of accident algorithms, and it is not analogous to the trolley problem that comes from the study of ethics. The authors of this paper contend that currently, only a human can make an acceptable conscious choice in this no-win scenario, where acceptability might differ from culture to culture. The authors are also cognizant of the risk of being caught in a prisoner's dilemma in pursuit of an artificial solution. An Artificial Creative Intelligence machine, developed for music improvisation, could provide the level of computational creativity required for making real-time choices to minimize physical, emotional, ethical, and societal damages.

Conclusion

Artificial Creative Intelligence requires a fundamental rethinking of how knowledge is managed and how intelligence is processed *in vivo*. The Insight-Knowledge Object model is offered as a new framework for scientists and engineers working in the fields of computer science, neuroscience, psychology, etc. Each level of knowledge object builds on the level beneath it. In order to continue up the hierarchy, knowledge must be captured and modeled differently from current and prior attempts at artificial creativity. In the IKO model, creativity is dependent on all levels of knowledge that come before it: void, null, datum, facts, know-how, memories, wisdom, patterns, predictions, and imitations. If any of these necessary levels are skipped, then the creative process collapses much like a house of cards, and Artificial Creative Intelligence remains out of reach.

Still unknown are the mathematics and implementations required to execute Artificial Creative Intelligence based on the proposed IKO architecture. The authors recognize blurry lines exist between levels that might direct future work in the direction of analog computing. To be determined is the exact nature of the insight functions. Are matrix algebra, the calculus, statistics, and probability sufficient to push the current AI boundary into creative insight processing? If not, what mathematics is needed for a software solution? How will quantum computing impact AI? Will an analog network of solid-state neurons play a role? Additionally, the biological basis for the IKO model has yet to be explored to understand the natural science foundation of the insight functions.

The issue of speed also requires solving. For ACI to be truly effective in numerous use cases, response times must match or improve upon that of human processing.

To be addressed as well is training an Artificial Creative Intelligence machine if deep neural networks are incorporated. Unsupervised training is particularly challenging in artificially judging pleasantness of newly generated music. Baker *et al.* (Baker, 2020) have launched an initiative in Human-Assisted Training-AI as a direction for future work in machine learning.

The above questions and issues require continued research and exploration across multiple disciplines including work in neuroscience, psychology, mathematics, computer science, electrical engineering, and philosophy.

The challenge is how to break the imitation barrier if at all possible.

Acknowledgements

In-kind support for this work was provided by Audio Intellegentia. The authors give thanks to the independent, anonymous reviewers of this paper for their feedback and suggestions.

References

- Abraham, A. (2018). *The neuroscience of creativity, Cambridge fundamentals of neuroscience in psychology*, Cambridge University Press, Cambridge, U.K.
- Aguiar, R., Liao, K. (2019). "Autonomous Haiku Generation", <https://arxiv.org/pdf/1906.08733.pdf>
- Andabi (2017). <https://soundcloud.com/andabi/sets/voice-style-transfer-to-kate-winslet-with-deep-neural-networks>
- Baker, J. K., Baker, B. J., Huang, X., Reddy, R., Mitchell, T., Garibay, I., Raj, B., Singh, R., Georgiopoulos, M. (2020). "An introduction to human-assisted training for artificial intelligence", <http://hat-ai.com/>
- Boden, M. (2004). *The creative mind: Myths and mechanisms* (2nd edition). Routledge, London, U.K.
- Chen, R. (2009), "Knowledge architecture", unpublished work.
- Chen, R. (2020), "Artificial inspiration machine", patent application filed with United States Patent & Trademark Office.
- Elliot, T.S. (1934). "The rock", Faber & Faber, London, U.K.
- Gao, Y., Singh, R., Raj, B. (2018). "Voice impersonation using generative adversarial networks", rXiv:1902.06840v1.
- Geravanchizadeh, M., Gavvani, S.B., (2020). "Selective auditory attention detection based on effective connectivity by single-trial EEG", *Journal of Neural Engineering*, 17 (2).
- IBM (2019). "The quest for AI creativity", <https://www.ibm.com/watson/advantage-reports/future-of-artificial-intelligence/ai-creativity.html>.
- Internet Archive (2019). <https://archive.org/details/GratefulDead>
- Kurosawa, A., (1954). "Seven Samurai", Toho Co., Ltd., Tokyo, Japan.
- Nyholm, S., Smids, J. (2016). "The ethics of accident-algorithms for self-driving cars: an applied trolley problem?", *Ethical Theory and Moral Practice* 19, 1275–1289.
- Olenik, A., (2019). "What are neural networks *not* good at? On artificial creativity", *Big Data & Society*, <https://journals.sagepub.com/doi/10.1177/2053951719839433>.
- Raj, B., Agarwal, S., Raj, T. J., (2019). "Music generation by deep learning", unpublished work from Carnegie Mellon University Language Technologies Institute.
- Rowley, J. (2006). "The Wisdom Hierarchy: Representations of the DIKW Hierarchy", *Journal of Information Science*.
- Runcon, M. A., Jaeger, G. J. (2012). "The standard definition of creativity", *Creative Research Journal*, 24(1).
- Sharma, N., Google (2008). "The origin of data information knowledge wisdom (DIKW) hierarchy", *Research Gate*. <https://www.researchgate.net/publication/292335202>
- Stein, M. I. (1953). "Creativity and culture", *Journal of Psychology*, 36(2).
- Turing, A. M. (1950). "Computing machinery and intelligence", *Mind* 49.
- Zeleny, M. (2005). "Human systems management: integrating knowledge, management, and systems", *World Scientific*.
- Zhou, Z., Yang, Y., Cai, Z., Yang, Y., Lin, L. (2019), "Combined Layer GAN for Image Style Transfer", *IEEE Xplore*.
- Zimmerman, K. (2017). "History of computers: a brief timeline", *Live Science*, <https://www.livescience.com/20718-computer-history.html>

*Where is the Life the authors have lost in living?
Where is the wisdom the authors have lost in knowledge?
Where is the knowledge the authors have lost in information?*
T.S. Eliot, 1934. "The Rock"