



DESIGN
THINKING
SERIES

Regenerating Learning

Transforming How You Learn
with Generative AI

Patrick Parra Pennefather

Apress®

Design Thinking

Design Thinking is a set of strategic and creative processes and principles used in the planning and creation of products and solutions to human-centered design problems.

With design and innovation being two key driving principles, this series focuses on, but is not limited to, the following areas and topics:

- User Interface (UI) and User Experience (UX) Design
- Psychology of Design
- Human-Computer Interaction (HCI)
- Ergonomic Design
- Product Development and Management
- Virtual and Mixed Reality (VR/XR)
- User-Centered Built Environments and Smart Homes
- Accessibility, Sustainability and Environmental Design
- Learning and Instructional Design
- Strategy and best practices

This series publishes books aimed at designers, developers, storytellers and problem-solvers in industry to help them understand current developments and best practices at the cutting edge of creativity, to invent new paradigms and solutions, and challenge Creatives to push boundaries to design bigger and better than before.

More information about this series at
<https://www.apress.com/series/15933>

Patrick Parra Pennefather

Regenerating Learning

Transforming How You Learn with Generative AI

Apress®

Patrick Parra Pennefather

University of British Columbia, Vancouver, Canada

ISSN 2945-7793

e-ISSN 2945-7807

Design Thinking

ISBN 979-8-8688-1060-2

e-ISBN 979-8-8688-1061-9

<https://doi.org/10.1007/979-8-8688-1061-9>

© Patrick Parra Pennefather 2024

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein

or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Apress imprint is published by the registered company APress Media, LLC, part of Springer Nature.

The registered company address is: 1 New York Plaza, New York, NY 10004, U.S.A.

About the Images in This Book

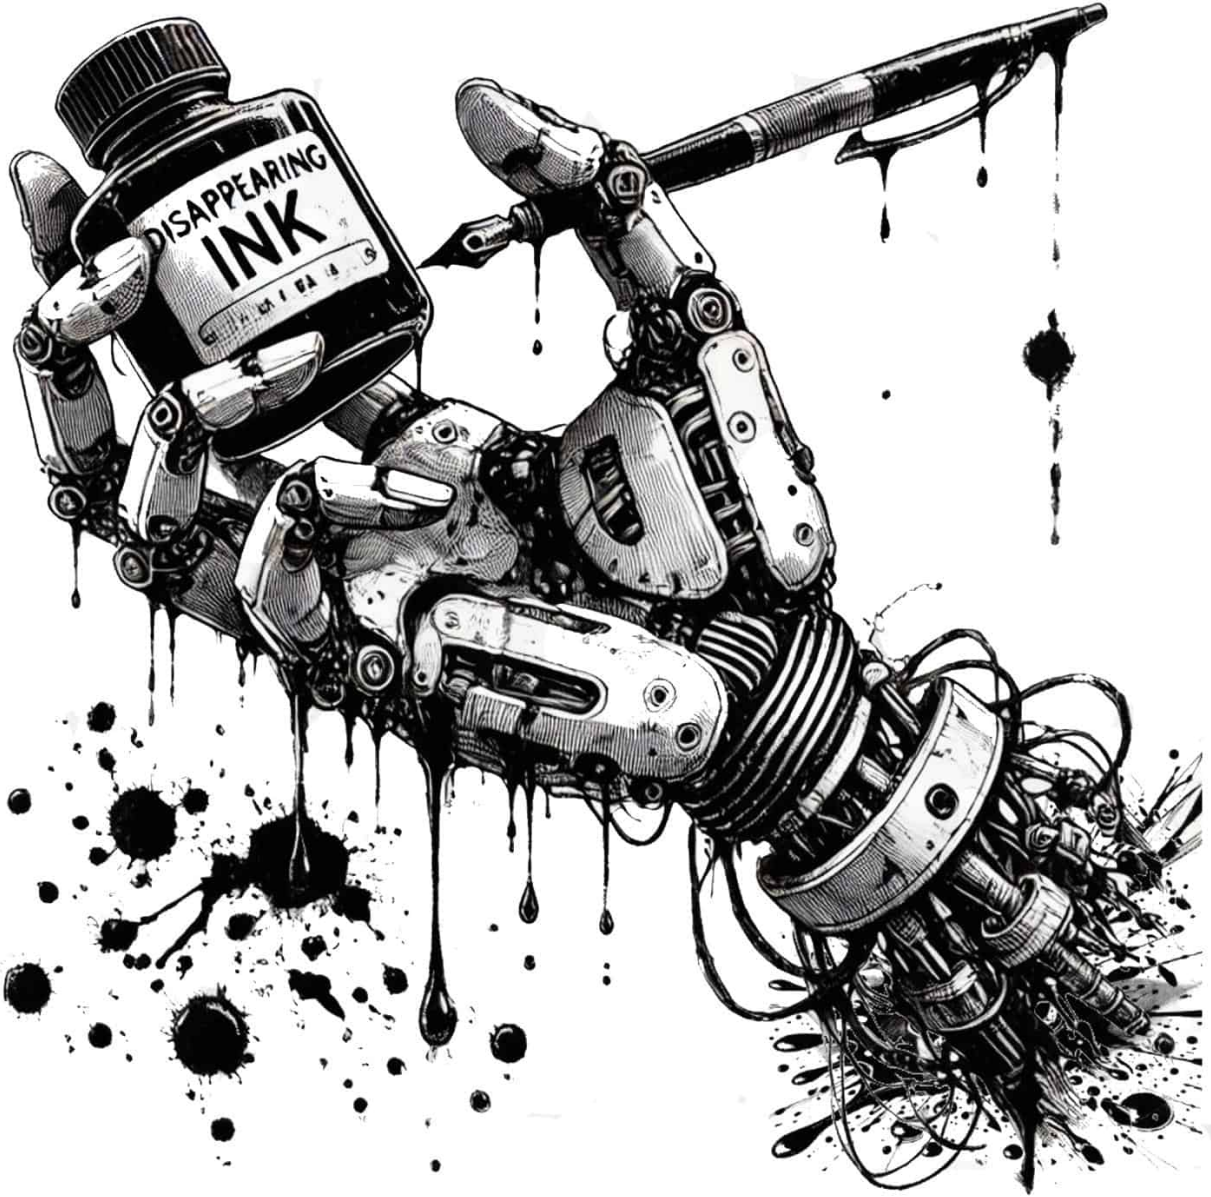


Figure 1 AI systems deal with obsolescence as much as ink has had to. AI-generated image

The human-machined images in this book have been extensively edited, using a variety of techniques and workflows, all of which were iteratively denoised by hand-me-down machine learning models on a private computer, not accessing the Internet, and by the time of this printing are likely obsolete. The workflows for the images have been as follows.

First versions of images were all generated using a customized and private machine learning model using OpenClip via ComfyUI on a PC, trained on an LAION-5B dataset.¹ All generated images were birthed from a text prompt and image owned by the author. In addition, some public domain images that were not part of the dataset were sourced and uploaded into the ML model from various repositories² focusing on the Xie Yi style of hand brush painting exemplified by Ming Dynasty (1368–1644) artist, Xu Wei—best known for his Xie Yi, or freestyle flower-and-bird paintings. Grapes (1521), by Xu Wei, was accessed via the Palace Museum, Beijing China. Other images that had stylistic features, which inspired the images in this book, were accessed via the Art Institute of Chicago,³ and all images included a CCO Public Domain Assignment. These include

- Kano Motonobu, “Ink Landscape,” hanging scroll, ink on paper, dated roughly 1499–1599 AD
- Xugu, Black Birds, hanging scroll, ink and colors on paper, 1824–1896
- Zhao Mengfu, Horse and Groom in the Wind, in James Cahill, *Ge jiang shan se – Hills Beyond A River: Chinese Painting of the Yuan Dynasty*, 1279–1368

Source and Accessibility: LAION collects a significant portion of its data from public domain sources and freely available web content. Public domain images are images that are free from copyright restrictions, allowing anyone to use, modify, and distribute them without legal repercussions.

Use of Datasets: Public domain images are often included in LAION’s datasets because they can be legally used to train AI models without the need for specific permissions. This accessibility is important for creating large

datasets necessary for training models to understand and generate visual content.

Contribution and Permissions: LAION also encourages users to submit images to their datasets. When users submit images, they typically grant LAION permission to include these images in their datasets. These contributions help expand the diversity and quality of the datasets, which is necessary for training AI models that need to understand a wide range of visual concepts.

Ethical Considerations: LAION takes care to ensure that user-submitted images are used ethically. This includes obtaining explicit consent from contributors and making sure that the images do not violate any privacy or copyright laws.

Quality Control: LAION's datasets are curated and often annotated with metadata that helps AI models learn from the images more effectively. This might include labels describing the content of the images, which are needed for supervised learning.

Open Source Ethos: LAION's datasets are generally made available under open source licenses, meaning that they can be freely accessed, used, and modified by researchers, developers, and organizations. This openness is a key part of LAION's mission to democratize AI research and development.

As mentioned in the book, many companies and organizations training generative AI go through various iterations after releasing their prototypes to the public. This is no different than those organizations dedicated to providing a corpus of public domain works for some of those machine learning models to train with. While no not-safe-for-work (NSFW) images were generated using LAION 5B in my own generated image sets, their recent update and the announcement inspired me to run all image prompts through their new downloadable corpus which cleaned their current set of links going to any suspected Child Sexual Abuse Material (CSAM) content.

All images were then highly edited in Affinity, removing backgrounds and unwanted figures, converting to black and white, adding single colors to some images as additional layers, and adding paint splattering effects ("ink and splatter" brushes) to strive for a visual consistency throughout the images used.

Some parts of each image were also removed in the process, and over a hundred black and white ink splatter content was generated and added as additional layers to all images. Screenshots of interactions with LLMs and visual models were all created by the author using a combination of Miro, Claude 3.5, Napkin AI, and ChatGPT4o.

Prompt Constraints

The following were eliminated from all prompts:

- Use of the phrase “In the style of,” followed by any living artist’s name or the name of any artist whose work is copyrighted
 - Use of the name of any living artist
 - Use of any artist who published work that is not in the public domain
 - Any reference to art from the 20th and 21st centuries
 - Any art created in the last 100 years; any photos of any art that were taken by living photographers for the past 100 years and/or not deemed public domain
-

Reverse Image Checking and Generative AI Image Identifiers

Despite the guarantees of “no copyrighted works” that LAION or any other generative AI system claims, to ensure all images were not sourcing from any living artist or copyrighted work, a number of approaches were taken. First, reverse image checks were implemented for all images using Google as it has the largest dataset of images. None of the final images used in the book were discovered using Google image checker. In addition, a total of 250 images that had any degree of likeness (according to Google) were scanned and cross-

checked for each image used in the book. Images or parts of an image that resembled a copywritten work were not included in the book and deleted from my hard drive.

All of the images were run through a minimum of two public generative AI image detector machine learning models. To be transparent, they were first tested with a real photo taken by the author to gauge the accuracy of their algorithms. Detectors also rely on a corpus of data but also use proprietary algorithms to detect if an image has been created with generative AI. It's good to test these systems as you would test any other generative AI platform.

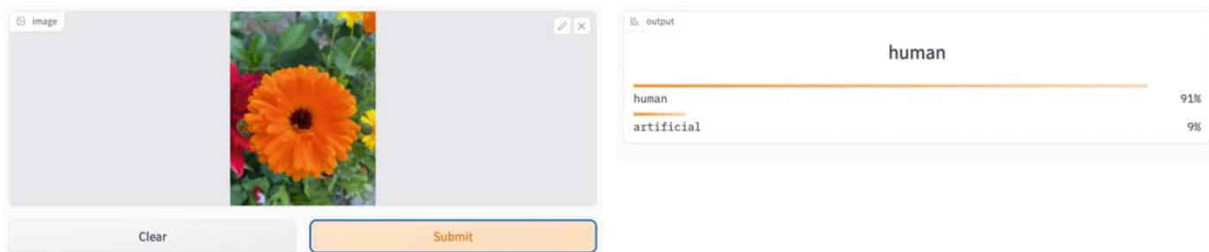


Figure 2 Original photo by author proved to be 91% human using Hugging Face's AI Image Detector⁴

Images in the book were tested with a minimum of two of the top-rated detectors currently available. This was done to better understand whether or not different machine learning models could detect if the images uploaded were AI generated.

Some provide more detailed analytics like FotoForensics and AI or Not (with paid-for advanced features), including the ability to predict if it's a Midjourney, Stable Diffusion, GAN, or Dall-E 3 image.⁵ Because of the level of post-generated editing, scores in Figures 2 and 3 were common, even though some images were drastically altered by me.

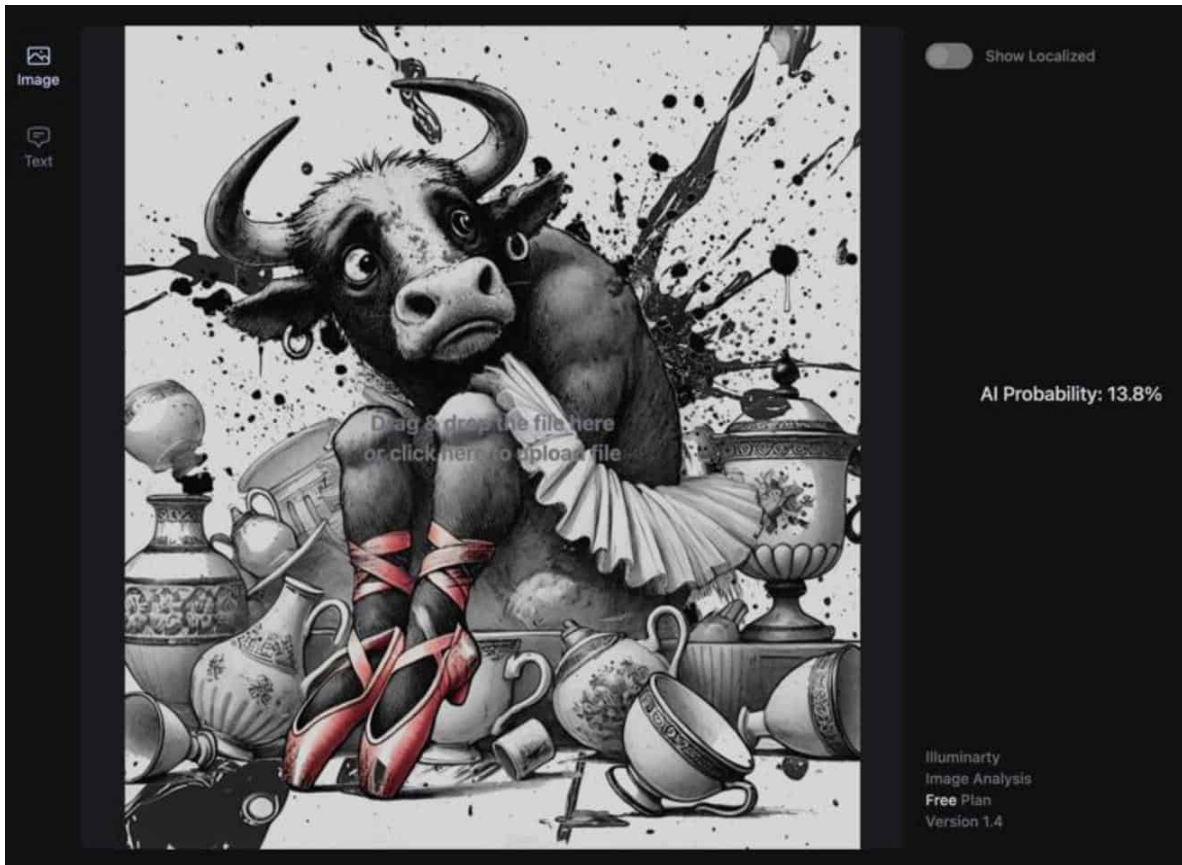


Figure 3 Bull in a China Shop tested with Illuminarty revealed a low probability of 13.8%. AI-generated image

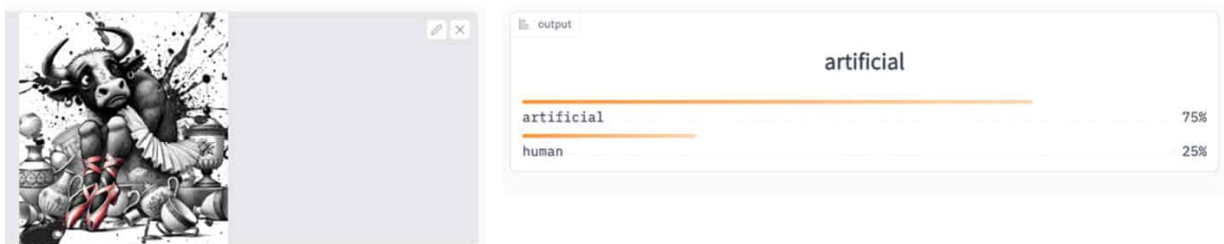


Figure 4 The same image tested with Hugging Face's AI Image Detector is likely more accurate with a 75% made with AI probability. AI-generated image

Human vs AI Probability



Figure 5 The same image used with Content at Scale AI Image Detection revealed a close similarity to Hugging Face's detector at 74% human created

Preface

Welcome to the brand-new wake-up call that is generative AI. For some, the technology is right by their bedside jolting them from peaceful sleep at 5am, provoking them to take action, learn about it, and unravel the mystery of how to integrate it within the work that they do. For others, it is a faint alarm sound heard at a distance, not even in their home, like a car alarm audible from two blocks away. Some more established AI researchers, developers, creatives, and companies are not aware of any alarm as they have been working creatively with the technology for a long time. Early developers of the technology, however, are sounding the alarm now in Oppenheimeresque tones and rhythms. No matter what you hear or the headline that captures your attention for about 12 seconds of your busy day, the annoying sound is persistent and is not really going away anytime soon no matter how much you may try and ignore it, take a break from it, or block it out. In the hurricane of alarm and adoption, this technology is inadvertently proposing a perfect storm of learning around its edges. In the eye of the hurricane, development teams of all sizes and value are busy in the labs crafting tools that they believe might solve human problems of all kinds. Extending the metaphor, as we move outward, we have creatives and organizations rethinking how this technology might solve a human problem, challenge established ways of learning and working, or provoke new forms of creation and artistic expression. As the hurricane hits our situational contexts, our work, and our homes, the tech blows around the messy and very human phenomena, like bias and lies, that we already know exist. That's OK because we are protected from it, in an underground shelter we have made for the big day. The perfect storm of learning provoked by generative AI is not just about learning how to use the technology to change human patterns of work and life. The technologies are reorienting how we think we learn, what we learn, what we need to learn, when and where we learn about knowledge production, how humans communicate with each other, and the

economic, social, political, creative, ethical, and technological factors that inform how we navigate human-influenced existence on this planet.

Narrow AI

In case you don't already know, generative AI is not that new of a technology, but it is an evolution of the development of machine learning models or systems that have been developed for over 60 years. Its development is also tied to the development of the computer. Generative AI that you might have played with, like ChatGPT or Midjourney, are considered narrow AI which have been designed for specific tasks like generating text or images. These dominantly open source systems have recently received neural steroids and a GAN facelift and are increasingly being developed privately for investors and more public consumption. Generative AI platforms have experienced increased investment to make the tech more accessible and profitable, with global use estimates of 250 million users in 2023, doubling from its use in 2020 and continuing to be adopted.⁶ For some sobering context to balance the hype or fear, that's less than 5% of a global population of about 8 billion estimated back in 2020.

That doesn't mean we are not affected by other narrow AI. As of 2024, there are approximately 5.16 billion active social media users worldwide, making up about 59% of the global population. Facebook remains the largest social media platform with over 3.15 billion monthly active users, representing approximately 39% of the global population. YouTube is the second most popular platform with over 2.5 billion active users, accounting for about 31% of the global population. Instagram is the fourth largest platform with more than 2 billion users, which is roughly 25% of the world's population. TikTok has rapidly grown to surpass 1 billion active users, representing around 13% of the global population.

While your mind has likely drifted with all the stats, this is important information to know. Why? All of these platforms leverage AI algorithms to make recommendations of other consumable content based on your interactions within each platform and with other humans (or cats) on those platforms. At times, user patterns are also analyzed to inform development teams as to what new features they might develop, and in what has become the norm, to target users with specific advertising content that the companies can profit from. The excellent book *How to Build Your Social Media Policy*⁷ (translated from the French) is worth a read for those who want to go deeper. As will be discussed in Chapter 12, many humans are also susceptible to algorithmic collusion, a new era of price fixing indirectly facilitated by third-party companies with industries like real estate and gasoline. Poetically, and in homage to David Bowie, algorithms are putting out the fire with gasoline.

Like other technologies before it, some humans in privileged contexts with sufficient means to access generative AI platforms are being asked to integrate it in our work, teaching and learning, reconcile its use individually based on many ethical dilemmas that surface from that use, and magically make our work more “efficient.” Is this technology a little bit different than other technologies that have promised to make human work easier, better, or automated? Who should be worried about being replaced in the work we currently do, when you have CEOs, wannabe thought leaders, futurists, and social media influencers constantly posting about the coming of even “smarter” and more “human” AI (AGI)? What do we need to know in case of its adoption in our work environments? Do we need to pay closer attention to that faint alarm sound? What do we do with all these statistics meant to forecast adoption while selling promises?

What have you learned about AI and where has that knowledge come from?

Artificial Puffery

You've heard the hype. AI is going to replace jobs. AI is revolutionizing how we work, making everything more efficient. It is surpassing human "intelligence" and will eventually have the capacity to act independently from a human operator. In case you didn't know, narrow AI systems have been spying on you for a while now. Not in the traditional surveillance tracking methods already established in certain city centers internationally, but in a more inconspicuous and annoying way, observing your search engine patterns and targeting you with ads accordingly. Public perception influenced by media portals has, over time, and thanks to box office hits like *Terminator* and others, persuaded many humans that AI is a technology that will also become sentient, control us, and take over all human work because all our tasks, creativity, and efforts can be automated. Beware of the large language models. They'll take your job. Beware of generative AI, as it plunders and extracts all knowledge, with or without permission or compensation. Worse is the seemingly incessant puffery, those promotional statements and exaggerations that make broad, boastful, or subjective claims about a sentient AI without providing factual evidence. Puffery is often used in advertising and marketing, but can appear in various forms of communication as hyperbole, propaganda, and sensationalism. Vague predictions are difficult to prove or disprove. Vague predictions are the lifeblood of spreading the use of generative AI. Let's see how an LLM deals with that last statement when prompted (or in the case below, provoked) to be vague in its own description.

VagueBot: Vague predictions sort of serve as a kind of cornerstone of AI because they perchance allow for a wide range of interpretations and polyphonic adaptability in various contexts. By not being overly unambiguous, these predictions can accommodate a multitude of unpredictable outcomes, making AI systems appear sort of more flexible and maybe even insightful. This broad applicability helps AI maintain a suspect amount of relevance across

diverse scenarios, more or less guaranteeing that its outputs can be seen as somewhat useful or occasionally accurate, regardless of the specifics. Consequently, the inherent ambiguity in such predictions contributes to the perception of AI as kind of a powerful and versatile tool, capable of maybe providing randomly valuable insights in an array of situations. So, in the future, AI might kind of continue to sort of evolve in ways that could be seen as pretty significant or maybe just kind of interesting or at the very least not trivial.

Misnomers, Context, and Semantics

For some more tangible background, the term machine learning emerged in the 1950s. It was introduced by Arthur Samuel, an American developer in computer gaming and artificial intelligence. Samuel is credited with defining machine learning in 1959 as a field of study that enables computers to learn without being explicitly programmed. On the surface, we can make sense as to what that means. We might translate Samuel's thoughts and likely summarize him incompletely by concluding that a computer learns without being programmed. If you start to repeat that sentence over and over, then go research if that's possible, and deepen your understanding of code, you will realize the irrationality of that statement. More specifically, a developer might say, "the software application is programmed to complete tasks and use logic to complete some other tasks without having to be always prompted to take on those tasks." You might even rip a coder away from their computer screen for 60 seconds and ask how they might actually say it: "the software is provided a means through which to determine (and replicate) patterns provided in the input, without explicit logic operating over the input domain (patterns often completely incomprehensible for humans to model)."⁸ If the very idea of intelligence regarding AI makes you feel stupid, you're not alone; at least you don't have to pass the Turing Test. Wrapping our heads around the language of how this technology actually works is an important step toward its broader

acceptance and usefulness. Tasking our LLM VagueBot to define intelligence when it comes to AI leads us to an interesting proclamation to take to heart.

VagueBot: The thing about intelligence in AI is that it's not entirely clear-cut; it sort of varies depending on the context and the specific applications. Sometimes, it's about being able to kind of solve problems or maybe recognize patterns, while other times it might be about sort of interacting with humans in a way that feels kind of natural or intuitive, I guess. So, overall, intelligence in AI is this nebulous concept that can mean different things in different situations, and it's all sort of about how we perceive the capabilities and behaviors of these systems in a kind of broad and maybe ever-evolving way.

AI or the debatable misnomer, artificial intelligence, is an umbrella term under which LLMs like ChatGPT and other probabilistic automated machines fall under. Machine learning models have been called intelligent because of their ability to perform complex tasks that would typically require human-like analysis and decision-making. The nuance is in how different people with different worldviews define "human-like." The word intelligence is also a very human-centric one that places its intelligence above any other sentient being's intelligence. AI has the word "intelligence" embedded within it, because of the correlation that it mimics human cognitive processes by analyzing large amounts of data, identifies patterns using statistical methods like probability (estimating the likelihood of events) and regression (finding relationships between different pieces of information), and creates new content or making guesses about future trends, behaviors, or outcomes based on those patterns. This ability to "learn" from data and generate results that appear to be based on understanding is why it's often referred to as "intelligent."

The seeming simplicity espoused by Alan Turing's seminal paper, "Computing Machinery and Intelligence"⁹ (1950), a foundational text in the field of artificial intelligence, will trigger your inner semiotician. Turing's approach to defining intelligence is referred to as pragmatic, emphasizing that instead of attempting to pin down a precise definition of intelligence, we should consider whether a machine can perform tasks or exhibit behaviors that would

be considered intelligent if performed by a human. This operational definition sidesteps philosophical debates about the nature of intelligence and focuses on practical criteria that can be tested and observed. Why the sidestep shuffle? Turing was not interested in philosophizing about the nature of intelligence likely to avoid controversy, in an attempt to advance the field, to establish testable criterion for machine intelligence, to avoid highly abstract debates, and likely to not slow down the progress of his work. Historically, the many debates on what constitutes machine intelligence have continued to this day.

“Can machines think?”, asked Turing. What is your question?

Despite the logic that AI is difficult to equate to human intelligence, as AI does not possess consciousness, non-simulated self-awareness, human creativity, emotional or spiritual intelligence, to combat the misunderstanding that gets communicated about the technology by tech companies with intentions that cannot be separated from capitalist structures, we need to come up with our own definitions of intelligence. Doing so will help counter the sensationalism around AI “replacing,” “outperforming,” and making decisions better than humans can.¹⁰ Disregarding the artificial part of the acronym, which has become more of a habit than you might think, is another risk in falsely comparing machine intelligence to be the equivalent of human intelligence. Can some narrow AI outperform humans in data analysis? Yes, foundational machine learning models can identify patterns in unstructured data very quickly if they are in a specific readable format and therefore will compute faster than the median human can. Capturing outliers or data that doesn’t fit a pattern the AI looks for, however, is a very human task. Reasoning what the implications of the data that we analyze are, and having that inform what actions to take based on that analysis, is also dependent on the features that define us differently, and less precisely, as differently human.

The Habit of Anthropomorphizing

Humans are fantastic at anthropomorphizing though. The tendency to anthropomorphize AI, or imbue it with human characteristics, stems from several factors. Research shows that when machines behave in ways that resemble human behavior, users feel more comfortable and connected. Assigning human characteristics to AI makes complex technology easier to understand. Anthropomorphizing AI can create an emotional bond between the user and the machine, enhancing user experience and satisfaction. This is especially evident in AI companions or virtual assistants, where human-like traits might engender trust and empathy.¹¹ Popular culture, through science fiction and media, has long portrayed AI as human-like entities. All these narratives, myths, memes, and tropes shape public perception and expectations, making it almost presumptive to attribute human characteristics to AI. Designers and researchers who imbue human traits onto AI have inherited old misnomers. Researchers often use terms typically associated with human skills and capacities when referring to AIs, emphasizing the supposed similarities between humans and machines. In the witty repartee that is “Artificial Intelligence Meets Natural Stupidity,” McDermott points out the misuse of words, even labeling them “wishful mnemonics” by researchers and programmers. Words now taken for granted, like “learning” and “intelligence,” were established in 1976, where AI emerged as a more robust field of practice.¹²

How do we explain the tendency to humanize AI among experts who should be well aware of AI's non-human nature?

Arleen Salles and others attempt to answer that exact question in their more recent article.¹³ The research and design of machine learning models has evolved, fragmented, and spun off into hundreds of different directions and applications, so it's easy to get confused by the language of learning and intelligence the inventors themselves persistently proclaim their systems to have.