Intuit Next Talk May 2, 2023 via Zoom

Implications of Generative Al for Software Product Design

Hugh Dubberly

"There is something legendary and historic happening in software engineering, right now..." "If you're not pant-peeingly excited and worried about this yet, well... you should be."

Steve Yegge, Cheating is All You Need, 2023
 https://about.sourcegraph.com/blog/cheating-is-all-you-need



Until recently, most AI programs were 'analytical'; they assess probabilities i.e., they recognize "things" and classify and count them.

These programs rely on a range of techniques from classical statistics to Bayesian inference.

AI =

- = NLP (algorithms for Natural Language Processing)
- = **CV** (algorithms for Computer Vision)
- = ML (Machine Learning: algorithms that can be trained)
- = **DL** (Deep Learning: ML using neural nets)

In the last few years, a series of 'generative' Al programs emerged.



Similar systems also generate protein structures and more.

Generating video



Midjourney



Synesthesia Studios



Runway

PART ONE

A bit of history for context

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023



The idea of creating beings has a long history in literature.



Talos, a 30m high automaton made to protect Crete, described in the *Argonautica* c. 300 BCE



Frankenstein, or the Modern Prometheus, by Mary Shelley, 1818



Rossum's Universal Robot (RUR), by Karel Capek, 1920

Plans for how AI might be made real go back about 75 years.

"What we want is *a machine that can learn from experience... the possibility of letting the machine alter its own instructions provides the mechanism for this.*"

— Alan Turing, Lecture, London, 1947.



Technical landmarks in the history of Al.

1936 Universal Turing Machine described, a foundation of computing		 1950 Turing Test (or Imitation Game), "Computing Machinery and Intelligence." 1950s Computational linguistics (or symbolic systems) attempted machine translation (Russian to English) by encoding morphology + syntax. 							2011 Siri speech recognition (2014, Alexa; 2016, Google Assistant) 2012 Large-scale image recognition, Hinton, Li, Dean, Ng	
		1956 Dartmouth Summer Research Project on AI, McCarthy + Minsky, with Newell, Shaw, and Simon.		1970 Minsky claimed, "from three to eight years we will have a machine with the general intelligence of an average human being."		1990 Second "Al winter" , funding again collapses, this time from Japanese government		pses, Generative adve Goodfellow	2014 Generative adversarial networks, Goodfellow	
		1958 Rosenplatt's perceptron learning algorithm, a neural network with only one layer for recognizing images (20x20 pixels)		1974 First "Al winter" , funding collapses, in part because computers were not powerful enough.		1998 Convolutional neural networks, LeCun		2017 Transformer arc	2017 Transformer architecture, Google	
1930	1940	1950	1960	1970	1980	1990	2000	2010	2020	
	1943 Neural Nets , McCulloch + Pitts, "A Logical Calculus of the Ideas Immanent in Nervous Activity."		1963 DARPA funded AI research at MIT. 1969 Minsky + Papert dismiss neural nets		1980 Expert systems, Feigenbaum 1985 Bayesian networks, Pearl			Dall-E 2 and ChatGPT, Op	2022 DenAl	
			,		1986 Back propagation train Hinton, Rumelhart, Will	ning, liams				

Al's history may be grossly simplified into three approaches.



2020s synthesis Combining both

logic regulating Bayesian networks Jiminy Cricket guiding Pinocchio

Sam Altman, et al.





One measure of Al's progress can be seen in game play:







LITERARY CHARACTER	BEATLES PEOPLE	OLYMPIC ODDITIES	NAME THE DECADE	
\$200	\$200	\$200	\$200	
\$400	\$400	\$400	\$400	
\$600	\$600	\$600	\$600	
\$800	\$800	\$800	\$800	
\$1000	\$1000	\$1000	\$1000	



Tic-tac-toe

Checkers

Chess

Jeopardy

1952

OXO developed at Cambridge was able to play perfect Tic-tac-toe.

1952

First checkers program, Arthur Samuel (the program learned)

1996

The **Chinook** team beats world's best Checkers player, wins US national tournament.

2007

The **Chinook** team announced that they had computed a formal proof that checkers was a draw.

1996

IBM's **Deep Blue** beats Chess champion Gary Kasparov for the first time.

1996

Deep Blue wins a whole match against Gary Kasparov.

2011

IBM's Watson won Jeopardy, beating two of the all time best players.

2015

AlphaGo beat Go champion Fan Hui.



ightarrow Texas hold 'em Go (Wei-chi)

Google's DeepMind program

2019

Pluribus (from CMU + Facebook) beat five other players in six-player Texas Hold 'em.

PART TWO

The economic drivers behind the current hype

11

The search-based advertising business is funding the current wave of AI.



\$488 billion

Advertising industry

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023



ByteDance



2022 advertising revenue

12

Microsoft has invested \$13 billion in OpenAl to reboot Bing.

OpenAI has made GPT. GPT is the Large Language Model (LLM) behind ChatGPT. Microsoft is building ChatGPT into its search engine Bing, in a project code-named Sydney.

Between 2019-2023, Microsoft invested \$3 billion in OpenAI, adding another \$10 billion in early 2023.

Google, Facebook, Baidu, et al. have similar efforts.





The launch of ChatGPT and Bing AI creates an existential crisis for Google, forcing both sides to escalate their technology as quickly as possible — a race to AGI.



improvements drive improvements in

At the center of all the hype is GPT (General Pre-trained Transformer) — one of a slew of Large Language Models (LLMs).

Backends	Frontends			
Large Language Models (LLMs)	Chatbots			
e.g.,	e.g.,			
GPT 4 (previously 2, 3, 3.5)	ChatGPT			
BERT	LaMDA			
Wu Dao	Zo, Tay			
LLaMA				
Image Models	Interfaces or editors			
e.g.,	e.g.,			
DALL-E	MSFT Designer			
Stable Diffusion	DreamStudio			
Midjourney	Discord			
	BackendsLarge Language Models (LLMs)e.g.,GPT 4 (previously 2, 3, 3.5)BERTWu DaoLLaMAImage Modelse.g.,DALL-EStable DiffusionMidjourney			

OpenAl's mission is to "build a platform that others are able to build businesses on top of." - Greg Brockman, President of OpenAl

LLMs are already being used as platforms by other developers:

- Bing + ChatGPT = Bing AI (AKA Sydney)
- Snapchat + ChatGPT = My AI in SnapChat Plus
- Notion + ChatGPT = Notion AI
- Quizlet + ChatGPT = Q-Chat
- Grammarly + LLM = GrammarlyGo
- Mem + ChatGPT = Mem, i.e., Mem was purpose built on ChatGPT ... (BTW, OpenAl invested \$23.5MM in Mem!)

Generative AI stacks are emerging.



Users

LLM's have become a new layer in the internet tech stack available on-demand via API.



The Long Tail theory was proposed by Chris Anderson in 2004 to describe the retailing strategy of offering more, less popular items. selling many copies of a few hits = selling a few copies of many titles



ltems

For example — a small, local bookstore will sell a few of the most popular items, and big chains like Barnes & Noble might offer a wider selection, while large online retailers like Amazon can offer a near-infinite selection of niche products.



Items

How might the Long Tail model apply to AI businesses?

Already, the race is on to commercialize LLMs; early efforts are focused on creating intelligent agents.



Domains

21

Likely "hits" are Al assistants in 'vertical' domains:



Domains

At the same time, more general-purpose 'horizontal' assistants are emerging, like ChatGPT, Bing AI, and others.



Domains

General purpose

ChatGPT

Bing AI

We also see other start-ups working on more general purpose Al assistants:

- For managing your calendar, to-do's, etc.
- For managing finances
- For managing your home, appliances, and car
- For gathering news

The emerging pattern is:

if your expert system already implements a knowledge graph, then augmenting it with an LLM-based interface is an obvious next step. Analysts suggest that this scenario will play out across the corporate world, as consultants (like Accenture) templatize this work and roll it out to their client base.

ChatGPT is already evolving into text-based intelligent agents the transition to ChatGPT + Cortana (Microsoft's version of Siri, Alexa, etc.) can't be far off.







The size of the LLM data sets (and their resulting models) has become huge. And they are growing exponentially, roughly 10x per year.

GPT 2 2019

1.5 billion parameters

175.0 billion

- **GPT 3** 2020
- **GPT 4** 2023

100,000.0 billion = 100 trillion

(other estimates project 170 trillion)



See also a similar chart in the Economist piece.

https://www.economist.com/interactive/briefing/2022/06/11/huge-foundation-models-are-turbo-charging-ai-progress

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

Sam Altman, CEO of OpenAI, suggests we might see a version of Moore's Law for AI.

•••



Sam Altman 📀 @sama

a new version of moore's law that could start soon:

the amount of intelligence in the universe doubles every 18 months

8:24 AM · Feb 26, 2023 · 3.7M Views

Moore's Law: The number of transistors on microchips doubles every two years Our World



Data source: Wikipedia (wikipedia.org/wiki/Transistor_count)

PART THREE

How does generative Al work?

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

GPT, LLMs, and all AI tend to work in a similar fashion: **Begin by compiling a large body of data.**



Then feed the data through an algorithm to create model — in this case a neural network with hundreds of billions of parameters.



Finally, send a prompt to the model, which will predict what comes next.



LLMs like GPT-3 generate text based on statistical models of language.

In other words, GPT-3 is essentially asking, "Based on the input text and what has already been said what is the most likely next word?"

'Parameters' are statistical relationships between pairs of tokens (e.g., words, phrases).

The model uses parameters to predict the next token. Through training, parameters are updated millions of times until the model has implicitly encoded language rules that humans are explicitly taught.



For text-to-image generators, the process starts in the text space the computation is done in the latent space and the output is in the pixel space.



Pixel space

image that humans can see







image decoder

output

user

PART FOUR

Now for the fun part, playing with ChatGPT

34

Entering this prompt,

"Write a biblical verse in the style of the King James bible explaining how to remove a peanut butter sandwich from a VCR."



Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

Entering this prompt,

"Write a biblical verse in the style of the King James bible explaining how to remove a peanut butter sandwich from a VCR."



ChatGPT responded,

"And it came to pass that a man was troubled by a peanut butter sandwich, for it had been placed within his VCR, and he knew not how to remove it.

And he cried out to the Lord, saying, 'Oh, Lord, how can I remove this sandwich from my VCR, for it is stuck fast and will not budge?'

And the Lord spoke unto him, saying, 'Fear not, my child, for I shall guide thy hand and show thee the way. Take thy butter knife, and carefully insert it between the sandwich and he VCR, and gently pry them apart. And with patience and perseverance, the sandwich shall be removed, and thy VCR shall be saved.'

And the man did as the Lord commanded, and lo and behold, the sandwich was removed from the VCR, and the man was saved.

And the Lord said, "Verily I say unto thee, seek not to put thy peanut butter sandwiches in thy VCR, for it is not a suitable place for such things. Rather, keep thy sandwiches in thy refrigerator or on thy plate, where they belong.

And the man heeded the Lord's words, and from that day forth, he kept his sandwiches in their proper place, and was saved from trouble and woe. Amen."

— https://www.nytimes.com/2022/12/05/technology/chatgpt-ai-twitter.html
The prompt is highly 'engineered':

The key request is to explain a task, a "how to " followed by a task definition " remove a sandwich from a VCR ," all of which is couched in a literary form: write "a biblical verse," and given a recognizable style: that of "the King James bible."



Juxtaposing the form/style with the subject adds to the fun, but we can't evaluate the result, other than it's surprising + amusing.

The prompt could be written in different ways, with each wording producing different results.



The same prompt can also generate infinite results,

because the processes start with randomly generated numbers, and they rely on probability.

Write a biblical verse in the style of the King James bible explaining how to remove a peanut butter sandwich from a VCR.

The prompt above also generated:

"Verily, thou shalt not place thy peanut butter sandwich in thine VCR, lest it become stuck and thou art unable to retrieve it. Instead, seek ye out a more suitable location for thy sandwich, such as a plate or a bag, and remove it from thine VCR with care. For it is written, "Thou shalt not put a stumbling block before the blind" (Leviticus 19:14). And the peanut butter sandwich in thine VCR is surely a stumbling block, causing confusion and trouble for thyself and others. So take it out, and do so with wisdom and caution."

And also:

"Verily I say unto thee, thou shalt take thine hand and graspeth the sandwich with great care, lest thou smiteth the delicate machinery of the VCR. Then, with gentle yet firm movements, thou shalt remove the sandwich and place it upon a table or other suitable surface. So shall thou avoid the wrath of damaged electronics and enjoy thy sandwich in peace. Amen."

You could also iterate on the prompt with different variations.

Prompt

Writea Shakespearean sonnetexplaining how toremove a peanut butter sandwich from a VCR

ChatGPT output

"Oh, woe is me! A sandwich stuck in there, Within my VCR, a peanut butter mess, How to remove it, I'm in deep despair, Without my tape player, I am distressed.

First, I must unplug it, and take care, For safety's sake, the wires I'll disconnect, Then open up the case, and breathe a prayer, Hoping that the sandwich won't infect.

With gentle hands, I'll try to ease it out, Without causing damage to the machine, Oh, may the sandwich give a little shout, And slip out smoothly, as if in a dream.

If all else fails, I'll take it to a pro, And hope that they can make my VCR go."

You could also iterate on the prompt with different variations.

Prompt

Write a limerick about

removing a peanut butter sandwich from a VCR

ChatGPT output

"There once was a sandwich so grand But it ended up in a strange land Stuck in a VCR It seemed quite bizarre Removing it took quite a hand!

With tweezers and patience and skill The sandwich was finally still Extracted with care It was worse for wear But the VCR was back in its thrill."

You could also iterate on the prompt with different variations.

Prompt

Write a haiku about

removing a peanut butter sandwich from a VCR

ChatGPT output

"Sticky sandwich stuck, VCR eats it hungrily, A buttery mess."

Prompt engineering (tinkering with prompts) is a new form of prototyping — 'prompt prototyping'.

"Prompting empowers a large, new, diverse user group (SWEs, UXers, PMs, ...) to rapidly prototype AI in minutes... requires little prior ML experience."

"As soon as you figure out the notion of prompting, you can start prototyping AI."

— Michael Terry, Google Research PAIR https://www.youtube.com/watch?v=pjE4IOqiNpA



What can be prototyped includes:

- Chat agent
- Classifier: Is this word a fruit?
- Information extraction: What fruits are mentioned in this shopping list
- Transformation: English-French; English-Python
- Interpretation/translation: Code synthesis
- Creativity support tools: WordCraft
- Social communities: Social Simulacra (Park UIST)
- Recommender systems: Food pairing
- Summarization
- Question answering

By the way, this list is a good summary of what LLMs seem to be good at.

'Prompt engineering', iterating on prompts to achieve different results, also applies to image generators.



Regal oil painting portrait of a tuxedo cat



Regal oil painting portrait of a tuxedo cat dressed as royalty



Oil painting portrait of a tuxed cat with black fur, bright yellow eyes, and white whiskers dressed in royal clothing



1700soil painting portraitof across-eyed, shorthair black catdressedlike royaltywith an order of the stick of the s

ChatGPT can even make simple diagrams and outline steps in processes like recipes.

Principles	
Design Principles	
Brand Guidelines	
Accessibility	
Cross-platform	
Performance	
Governance	
1	
+V	+
Components	
Typography	
Color Palette	
Icons	
Buttons	
Input Fields	
Tooltips	
Notifications	
Mavigation	
Cards	
Lists	
Tables	
Accordions	
Progress Indicator	
Badges & Tags	
*	
Implementation	
HTML/CSS/3S	
Component Library	
Design Templates	
Documentation	
Versioning	
Testing & Validati	on
Contribution	
Training & Educati	on]
1. Principles: This layer includ	les the foundational guidelines and policies that d
design system's creation a accessibility and ocusroad	nd usage, such as design principles, brand guidel
2. Components: This layer co	ue. ntains the visual and functional building blocks of
design system, such as typ	ography, color palette, buttons, input fields, and
components.	
 Implementation: This layer Evolution: Such as MTM, Inc. 	deats with the practical aspects of applying the d SUS code, component libraries, dealers torgetate
documentation untriorion	aus code, component iuranes, design templates,

sual language across the entire as

3		
	String Diagram for Carrot Risotto with Chile Crisp Recipe	
		1
		\
		1
	[Add Cheese & Butter]> [Step 6]> [Stir & Season]	
		> [Step 7]> [Se

lask	Start Time Dura	t10n
Prep Carrots	00:00 05:00	9
Heat Oven	05:00 05:00	9
Toss Carrots	10:00 05:00	9
Bake Carrots	15:00 15:00	э
Warm Stock	20:00 05:00	Э
Cook in Saucepan	25:00 07:00	9
Cook with Wine	32:00 03:00	э
Stir in Rice	35:00 02:00	9
Add Warmed Stock	37:00 15:00	9
Stir & Season	52:00 05:00	9
Serve	57:00 03:00	9
Total Duration	L 60:00	 a

The time format used in the chart is in minutes. The duration for each task is an estimation and can vary depending on your familiarity with the recipe and your cooking skills. The chart is meant to give you a general idea of the order and duration of each step in the recipe.



People are also using ChatGPT to generate code for rendering diagrams and charts.



Graphologue: Exploring LLM Responses with Interactive Diagrams Haijun Xia, UCSD, HCI

https://twitter.com/haijunxia/status/1646917869115166720?s=46&t=TjLBVtVf7X7T7IC_uxXAGw

See also: https://www.akkio.com/chat-explore https://www.youtube.com/watch?v=7v6u7m_x3ls https://medium.datadriveninvestor.com/chatgpt-helps-you-to-create-graphs-with-python-in-5-steps-55418100d5c1 https://github.com/varunshenoy/graphgpt

https://xmind.app/blog/chatgpt-and-xmind-how-to-create-a-mind-map-with-chatgpt/



They're not perfect

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

48

LLMs may "hallucinate" —

i.e., make-up stuff.

Some have called LLMs

- Stochastic parrots
- Bullshit generators
- A blurry JPEG of the web

https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web

"A raw LLM is like a Harvard CS grad who knows a lot about coding and took a magic mushroom about 4 hours ago, so it's mostly worn off, but not totally...

You get the LLM to draft some code for you that's 80% complete/correct.

You tweak the last 20% by hand."

— Steve Yegge, Cheating is All You Need, 2023 https://about.sourcegraph.com/blog/cheating-is-all-you-need

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023



For example, GPT-3.5 has hallucinated movies in which a character shoots off a flare gun.

Prompt

"Write a list of movies in which" a character shoots off a flare gun." ChatGPT output

- 1. The Hunger Games: Mockingjay Part 1 (2014)
- 2. The Revenant (2015)
- 3. Cast Away (2000)
- 4. *The Martian (2015)*
- 5. *Life of Pi (2012)*
- 6. *The Grey* (2011)
- 7. Captain Phillips (2013)
- 8. Jaws (1975)
- 9. Crimson Tide (1995)
- 10. The Abyss (1989)*"

* - In The Abyss, a character ignites a flare, but not a flare gun.

"Here is a list of movies in which a character shoots off a flare gun:"

Jaws is the only movie on this list that includes a flare gun.

Image generators 'hallucinate' as well, making errors in certain visual details.









Al systems should explain how they work, enabling people to understand their strengths and limitations and how we can use them best in different contexts.

The goal is to make it possible for humans to understand how and why an AI system makes certain decisions or produces certain outputs. The focus will shift to asking questions like "Why did the abstract generator emphasize that portion of my paper?"

Making the AI understandable will help to increase the trust that people have in AI systems, reduce the risk of bias (depending upon the training data used), and improve the overall performance of AI systems.

> LLM stands for Large Language Models. These are algorithms that can detect, summarize, translate and predict text sequences based on massive datasets ¹. They use natural language prompts which make them popular¹. LLM-based chatbots allow customers to interact with businesses in a more natural and conversational manner². Some examples of LLM chatbots include ChatGPT, Bard, Ernie and Claude 3.

Is there anything specific you would like to know about LLM chatbots?

Learn more: 1. techxplore.com 2. venturebeat.com 3. linkedin.com +2 more

Bing AI chat response annotated with citations

1 of 10 🔵

Organizations should implement AI systems responsibly, with governance structures in place to:

θľθ

Minimize bias in algorithms and data

Build responsibility into your AI to ensure that the algorithms – and underlying data – are as unbiased and representative as possible.



Ensure transparency into AI systems

To build trust among employees and customers, develop explainable AI that is transparent across processes and functions.



Protect privacy + security of data

Leverage a privacy and security-first approach to ensure personal and/or sensitive data is never used unethically.

https://www.accenture.com/us-en/services/applied-intelligence/ai-ethics-governance

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023



Create opportunities for individuals

Empower individuals in the organization to raise doubts or concerns with AI systems and effectively govern technology, without stifling innovation.

Generative Al in the design process

55

Now is the time for designers to jump in to ensure we are addressing people's and society's needs.

We need to clearly understand the implications when technology takes actions on behalf of humans.

Designer Mike Kuniavsky argues that much of what we call 'creativity' is also a sort of 'auto-complete'.

Designers start with a 'fuzzy' idea for something it's inchoate, vague, half-baked — low fidelity (lo-fi). Much of the work of designing is filling-in the details, which Kuniavsky claims is a kind of auto-complete. Generative AI systems can fill-in the details of a "prompt", returning a high-fidelity image or text almost instantly.



Sketching

Planning

Modeling

Building

Sydney Opera House Jørn Utzon

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023



The Bull Pablo Picasso

Human designers also go through a series of iterations —

the sequence of feedback loops that clarify context and intent, and fill-in the details on both sides of the equation (context and form), in a conversation with stakeholders and the situation.



Designers also 'back-track' when a particular path leads 'off-course' or the designers choose another goal (once they see where they ended up).



Fidelity of output may become a 'knob' on generative AI systems, though getting them to abstract may be quite a challenge for engineers.

Building 'trees' of increasingly higher fidelity results seem potentially quite useful.



Some generative AI tools include various controls for the user to tune the output.







Bing chat has recently implemented a 'conversation style' control.

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

Adobe Firefly also includes controls.







Generative AI will transform design and other fields.

Some tools have already emerged, with more to come.



Generative AI will be human amplifying vs. human replacing.

Partnering with people to play various roles:

- Inspiration engine
- Brainstorm partner
- Co-writer / first draft creator
- Content adapter
- Brand enforcer
- Research partner
- Asset enhancer / organizer
- Financial, legal, or medical advisor
- Life coach

. . .

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

For a software system used to make consulting proposals, an LLM can be used to automatically generate custom information for each client.



Some designers have begun to use generative AI throughout the process:

Research synthesis

Transcribe interviews with **Reduct** and highlight stand-out quotes.

Run the quotes through **Notion Al** with a prompt like: "What are 10 key takeaways from this conversation?"

Also, **Vowel**, for automated meeting summaries.

Concepts & Ideation

Use AI-based image generators like **DALL-E** or **Midjourney** for inspiration.

Also, **Fermat**, for concept creation.

Galileo Al or Genius for generating UIs.

Critique

Use **ChatGPT** etc with the prompt: "Give constructive feedback on the following ..."

Marketing Copy

Use **ChatGPT** to write a first draft of marketing copy, then edit.

https://www.ideo.com/journal/5-ways-were-using-ai-at-work



Decks

Use **Tome.app** for creating pitch decks or other presentations.

Designers have at least three roles to play in shaping the development and use of generative AI as a new technology:

Role 1:

Designing new ways to interact and understand AI

Role 2:

Bringing the designer's lens to business growth and differentiation with AI

Role 3:

Designing in governance, safety, and ethics by co-creating with the technologists

Role 1: designing new ways to interact and understand Al

Marshall McLuhan once said that "we shape our tools and our tools shape us".

We will need to redefine user experiences.

The right interface to make the experience compelling and simple, and the 'magic' to feel transformative without overpromising will be table stakes.

Role 2: Bringing a designer's eyes to business growth and differentiation with AI.

We can help our clients design their response to AI with a human lens by amplifying the company's purpose.

It will take a village —

we'll work with data scientists, AI experts, psychologists, learning experts, business strategists, designers and design researchers.

To mediate the collaboration of the team, a new view of the system as a boundary object will be needed.

Role 3: Designing in governance, safety, and ethics by co-creating with the technologists

These teams will have a great responsibility to design for good from the start with ethical design principles applied as the systems are conceived.

That means co-creating from the start to design for the intended consequences, while using foresight to continually understand the unintended consequences.

PART SEVEN

Some final thoughts

Dubberly Design Office + Accenture Song + Implications of Generative Alfor Software Product Design + May 2, 2023

71

Experts are both very skeptical and very concerned.



"Large language models have limited reliability, limited understanding, limited range, and hence need human supervision."

— Michael Osborne, Professor of Machine Learning in the Dept. of Engineering Science, University of Oxford, January 25, 2023

"We are in a once-in-a-lifetime moment: just as" the atomic bomb forced the world to create new governing structures such as the United Nations, the unfolding A.I. developments are global in scope and demand a collaborative paradigmatic response."

— Tristan Harris, Executive Director & Co-founder, Center for Humane Technology, February 27, 2023


In the short term, we tend to overestimate a new technology in the long term, we tend to underestimate it.



Hype Cycle for Artificial Intelligence, Gartner, 2022

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

How will Al affect designers?



"Everything you do as a profession is going to have a Copilot"

- Reid Hoffman, founder of LinkedIn

"AI will not replace you. A person using AI will."

- Kai-Fu Lee, Ph.D., AI expert, former president of Google China



In the late 1960s, Nicholas Negroponte imagined an "Architecture Machine," not computer-aided drafting, but computer-aided design.

"I examine the design process in terms of its being conducted (not necessarily by professionals) in concert with computers in particular, with a class of computers that someday may exhibit intelligent behavior. I consider three potentials of the computer:

- the computer as a **designer**,
- the computer as a **partner** to the novice with a self-interest, and 2
- the computer as a physical environment that knows me." 3

What Negroponte dreamed of was a computer that could converse.

"I believe that computers have the potential for assuring a responsiveness, individuality, and excitement in all aspects of living, to a degree hitherto unseen...

I am interested in the rather singular goal of making the built environment responsive to me and to you, individually, a right I consider as important as the right to good education."

– Nicholas Negroponte, Soft Architecture Machine, MIT Press, 1972



Building on Negroponte's dream of the Architecture Machine, we might propose "generative conversation" as a goal of "generative AI".

A generative conversation is a conversation between two colleagues, which builds on the knowledge and experience of both, in areas of importance and interest to both, in ways that generate new understanding, agreement, and delight.





APPENDIX

Resources

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

78

Play with Al

Dall-E image generator https://labs.openai.com/

ChatGPT text chatbot assistant https://chat.openai.com/

Bing AI search + text chatbot assistant https://www.bing.com/

Tome presentation generator <u>https://beta.tome.app/</u>

Copilot programming assistant https://github.com/features/copilot

Learn more about Al

Huge "foundation models" are turbo-charging Al progress June 11, 2022 https://www.economist.com/interactive/briefing/2022/06/11/huge-foundation-models-are-turbo-charging-ai-progress

ChatGPT Is a Blurry JPEG of the Web Ted Chiang, February 9, 2023 https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web

What Is ChatGPT Doing ... and Why Does It Work? Stephen Wolfram, February 14, 2023 https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/

Bing's A.I. Chat: 'I Want to Be Alive' Kevin Roose, February 17, 2023 https://www.nytimes.com/2023/02/16/technology/bing-chatbot-transcript.html

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

APPENDIX

AGI

Dubberly Design Office · Accenture Song · Implications of Generative Alfor Software Product Design · May 2, 2023

AI (Artificial Intelligence) is a catch-all term.

It raises questions:

- What is intelligence?
- And how might we recognize it?

Alan Turing suggested the "imitation game" now known as the Turing test. If a person cannot distinguish a machine from another person when interacting with both, then we might say the machine is "intelligent."

We might say such a machine exhibits Artificial General Intelligence (AGI).

Special thanks to Todd Elliott Roger Meike Toni Belknap Eric Knudtson

And acknowledgements to: Gavin Miller Ryan Reposar Jake Sheiner Ian Shadforth John Cain Shelley Evenson

Presentation posted at presentations.dubberly.com/generative_ai_implications.pdf