

FIELD



generative
AI


* NOTES *



Matt Wood


art by @abbiemhart





Generative AI is a journey - not a race - that is only just underway and will make explorers of us all.

A race is a competition. It has one winner. And many losers. There is no winner take all in generative AI, since the benefits are so broadly accessible.



We are all on a journey to understand, experiment, and safely apply these new

capabilities in AI to our products, experiences, and businesses.

THAT JOURNEY JUST STARTED.

Rare periods of discontinuous change



present rare periods of opportunity,

and the



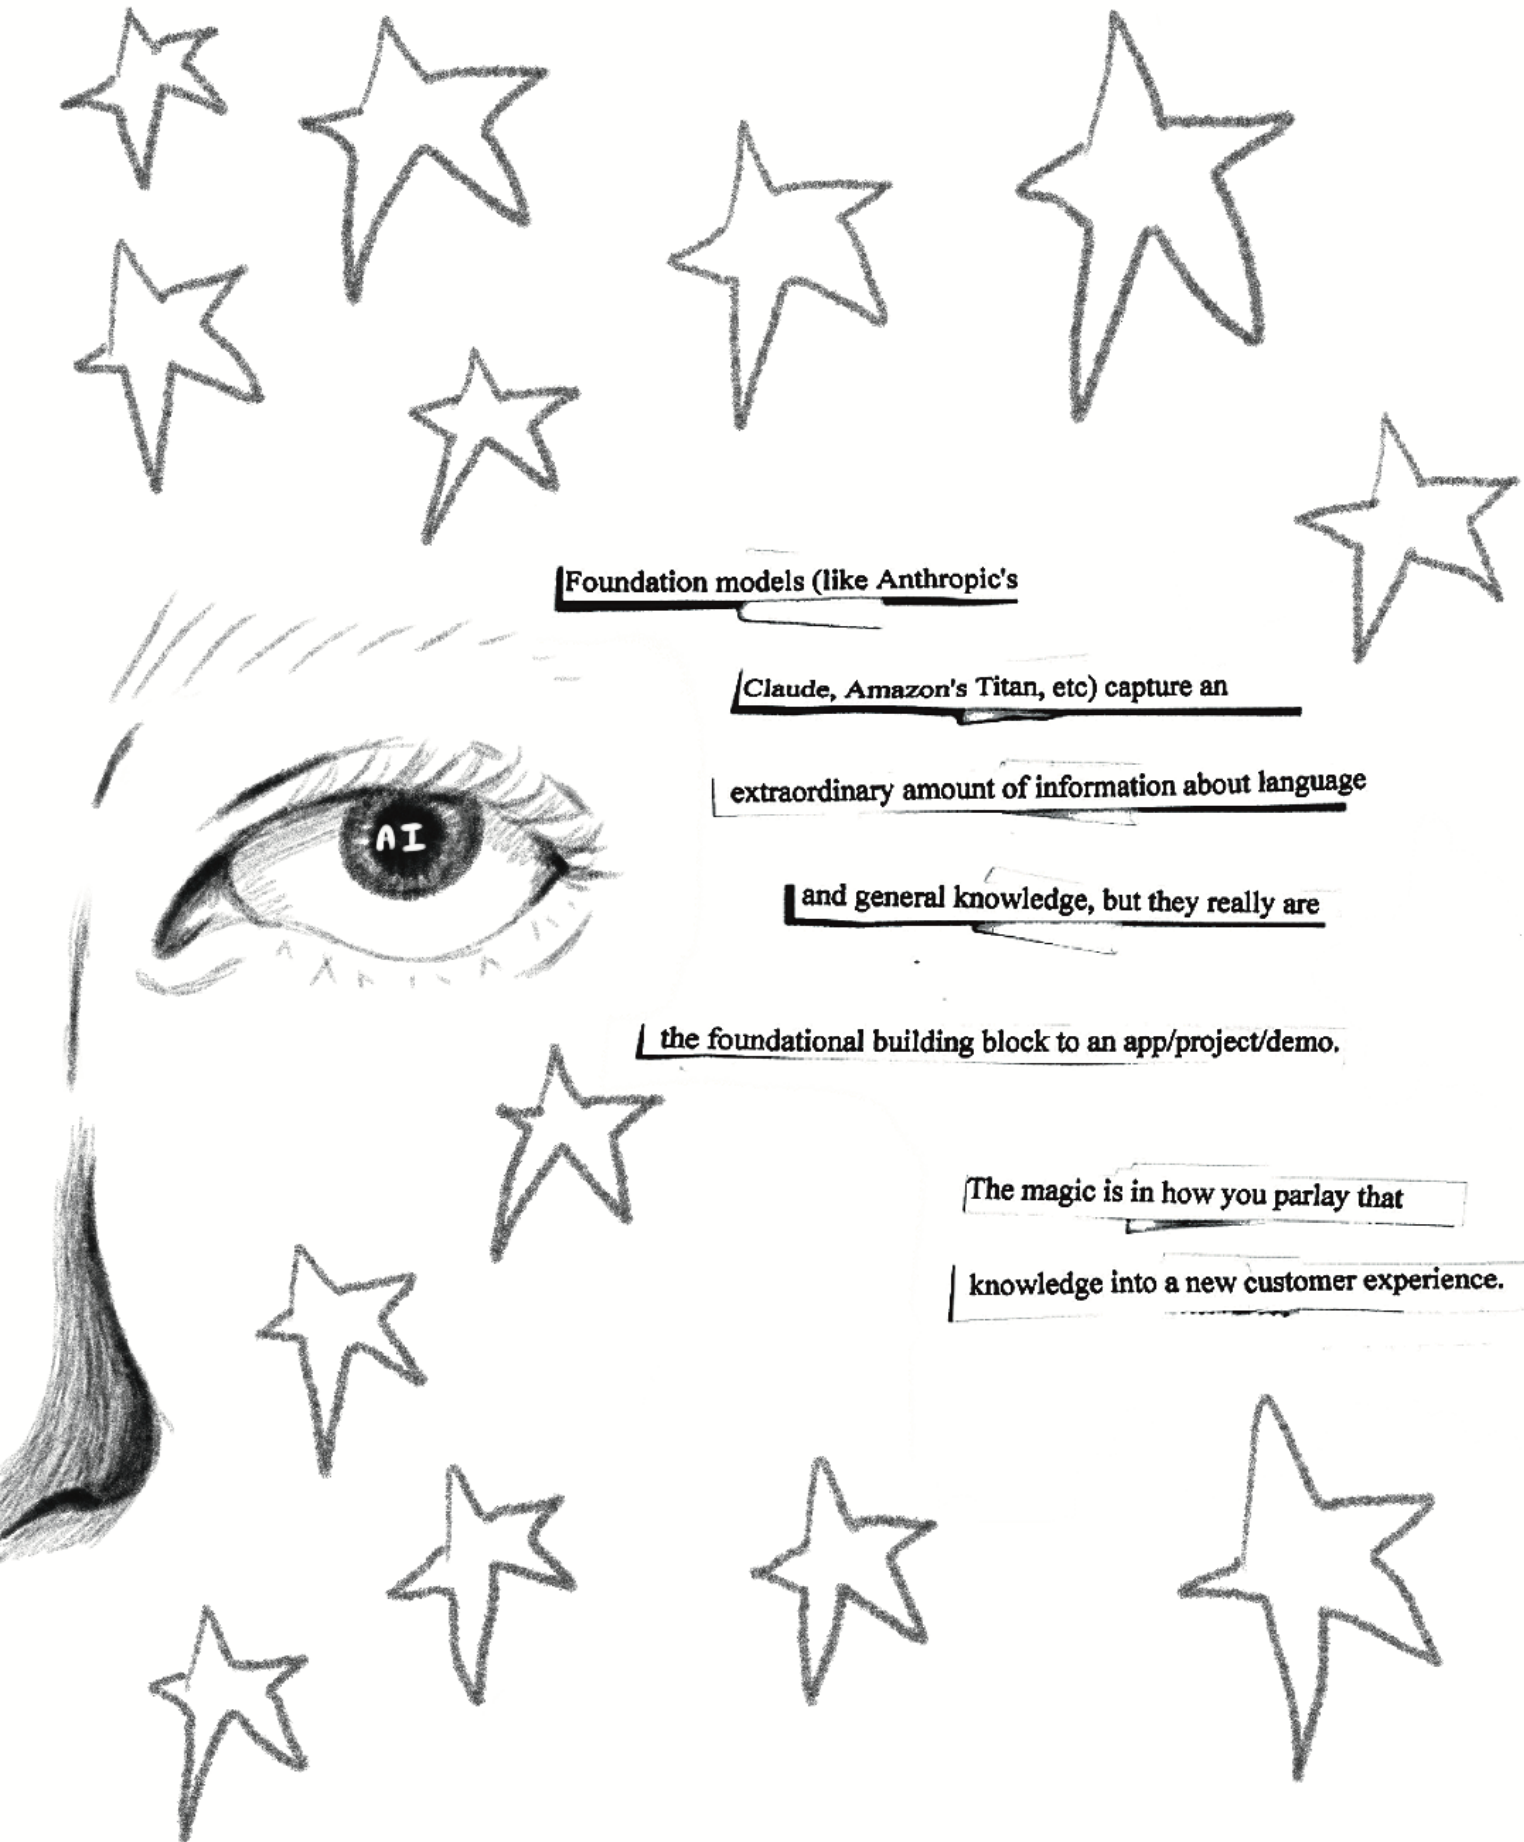
opportunity here for customer experience



invention, simplification, and innovation is truly



EXCITING



Foundation models (like Anthropic's

Claude, Amazon's Titan, etc) capture an

extraordinary amount of information about language

and general knowledge, but they really are

the foundational building block to an app/project/demo.

The magic is in how you parlay that

knowledge into a new customer experience.



|Knowing where and how customers will interact| with that experience can really help
when thinking through where and how you will want to build on top of the foundational
model with your own prompts, data, fine-tuned examples, and so forth.

EXPERIENCE

common use-case patterns of generative AI to help drive more efficient workflow:

1/ creative tasks (where your end user usually knows pretty much what they want, and so the generative model can give them a starting point or jumping off point)

2/ behind the scenes (where your end user won't even know they are using an LLM, but the experience will just be better because of it)



3/ question and answer (where a novice in a particular field wants to learn about a new topic, event, domain, and can do so through a series of q&as)

4/ expert decision making support (where experts collaborate with one or more AI agents in order to trouble shoot, goal seek, or solve complex problems)





Each interaction with an LLM provides three major ways of customizing or improving the output of the very general, very large, foundational model under the hood.



1/ engineering the input prompt (and context window)
- this is everything you provide to the model as an input
- the instruction, question, prompt, plus the extra additional supporting information. You can use this to customize the system at 'run time' for things like tuning writing style, referencing additional technical materials, data structures to model, narrative continuation, real time data, and so on.





2/ customizing the model - in some cases (like Titan),

you can continue pre-training the LLM with your own data

(great if you have large volumes of existing unstructured or


semistructured internal text data on your own business or use cases),

or fine-tune the model with structured data (using things like

question and answer pairs, giving examples of tone, personality,

or technical data)





3/ evaluating the response - where you response meet the needs of the user?" If so, great! model to reinforce a positive interaction (making you'll want to annotate and correct the output, and via improving the prompt, or improving the model.

will want to answer the question: "did this Grab that nugget and feed it back into the it more likely to reoccur next time). If not, feed the improvements back into the system

My advice (you may know a better way), is to focus on the experience, process, and steady progress along that journey, rather than an arbitrary, point in time destination or trying to call a winner.

TO QUOTE THE GREAT PHILOSOPHER
DRAKE:

"Sometimes it's the journey that teaches you

a lot about your destination."



FIELD



NOTES

