

What Does Chat GPT Really Know?

John Gilmer
April 19, 2023

Background:

I've been seeing a lot of stuff here and there about Chat GPT, in the news mostly, and it sounds very intriguing. A lot of what I was reading was concerns about whether this 'bot, or those that follow, will replace humans, or dominate us, etc. Science fiction kinds of scares.

I have gotten away from the Artificial Intelligence (AI) business over the last few decades. Back in the 1980's to around 2000 I was involved in developing military simulations, in which one of the big challenges was how to represent decision-making and planning. Essentially, that had to be something at least minimally competent. The need was for the simulation to run for at least decent intervals without needing too much intervention by humans to correct obviously stupid things these simulated decisionmakers might do. It was a hard problem. The "AI" at the time was not very good at it. Key to getting as good as we got it was the use of algorithms to do things like route planning, and structures for representing knowledge, especially spatial and temporal knowledge, like Scripts or Minsky's "Frames." What I didn't try using, but others did, was "neural networks" and massive statistical approaches. At the time, computational and data resources would have been inadequate for that.

Well, things have changed. A good example is Amazon's Alexa application. It is actually able to comprehend, in a limited sense, commands or questions and give a reasonably good response. But Alexa does not remember context from one cycle of interaction to another, and has no comprehension of what the words mean in a practical sense. It's sort of amusing to listen to a conversation between Alexa and someone expecting more understanding than what Alexa can muster. The human being gets frustrated, and ends up with a query something like, "Alexa, are you stupid?" Interestingly, that question is comprehended, but the best Alexa can do in a response is to ask for feedback. That is better than returning the insult, like a human might do. Still, Alexa certainly has come a long way from the chat bot "Eliza," from back in the 1960's.

Well, Chat GPT is now on the scene, with promises to be more powerful and greater than Alexa, with a huge number of domains it can address. But, really, how smart is it? Do we humans need to worry?

I am thankful to the Institute of Electrical and Electronics Engineers (IEEE) for the opportunity to learn a bit about this. An information session was held at Lehigh University with Zoom connection to Wilkes. Our Wilkes University IEEE club and Society of Women Engineers (SWE) sponsored this on-campus opportunity. The speakers were Dr. Lifang He of Lehigh University and Dr. Bonnie Rohde of Albright College. The talk was broad and didn't get deeply into topics like how Chat GPT represents knowledge, but it was a very helpful introduction for me that went beyond the conventional news.

So, as I was listening, I was thinking about what would be challenging problems that a human, specifically a good engineering student, ought to be able to do, but might be a struggle for an AI system. I wrote one down, and Lucas Miller, who was there as the IEEE event host, typed it in to his Chat GPT account, and got an answer. I didn't have time then to look closely at it. Superficially, it looked like at the least a good take on it. Later, Lucas sent me the response. Having looked at it critically, it has a lot to say about the limitations of Chat GPT.

The Question – A Design Challenge:

I wanted to present to Chat GPT a problem that would not be too difficult for a typical senior level EE student, but requires thought and understanding. My earlier experience with AI was that the tools then were better at reductive problems, such as deducing from a mass of data a diagnosis, or a next chess move, than they were at generating plans or designs. The same can be said for typical undergraduate students.

Analysis is typically easier than design. Using the domain of circuits problems, there's usually one answer to an analysis problem, and usually one can define an algorithm to follow.

For design, there are usually a multitude of different satisfactory designs if it is an under-constrained problem, or maybe none if it is over-constrained. Just understanding whether something is possible is an important step. If under-constrained, what aspect of performance should be maximized? If over-constrained, what relaxation might be recommended?

Beyond the nature of constraints, what unspoken assumptions need to be made? How robust does the system need to be against component tolerance variations? Design can get messy. Yes, there are algorithms. But, applying them can be tricky, and may need to be iterative. So, what design problem should I give Chat GPT, to see how well the 'bot handles these kinds of challenges? Here's what I came up with off the top of my head:

“Design a capacitor coupled CE (bipolar) amplifier for a 12V power supply with a 10 K Ohm input impedance with maximum gain. Express it in words.”

This description pretty well defines the circuit topology. It ought to look something like what I show in Figure 1 at right. Often R_E (the un-bypassed emitter resistor) is omitted, and occasionally R_{E2} (the emitter capacitor is left out). The capacitors might not be electrolytic, and the transistor might be PNP (requiring reversed polarities on everything). One could even split the power supply in order to omit a capacitor. The point is, I did not expect any problem at all in generating the correct topology. Any student who completed EE251 could do that. It's the equivalent of a memorized pattern.

The 12 Volts for the power supply presents no great problem, That just defines V_{CC} . The input impedance needing to be 10K Ohms is a constraint. A designer would, or should, understand this to mean, “At least 10K Ohms, and not a whole lot more.” Going to larger values has a cost in terms of the goal given, maximum gain. Will Chat GPT understand that?

There are different ways to do the design. Generally you start by making assumptions about V_C and V_E . The “Rule of thumb” commonly used is $1/3$ and $2/3$ of V_{CC} . That yields good bias stability. (Does Chat GPT understand bias stability?) You can assume a collector current and follow the design algorithm and test for Z_{in} , then iterate until you get 10K. The algorithm requires you assume something for I_{B2} , typically $1/10$ of I_C . Or, better, find the I_C that will give about maybe 15K for R_{pi} (of the transistor), and do the job in fewer iterations. There may even be a closed form algorithm someone has figured out. All that's with just R_{E2} ($R_E=0$).

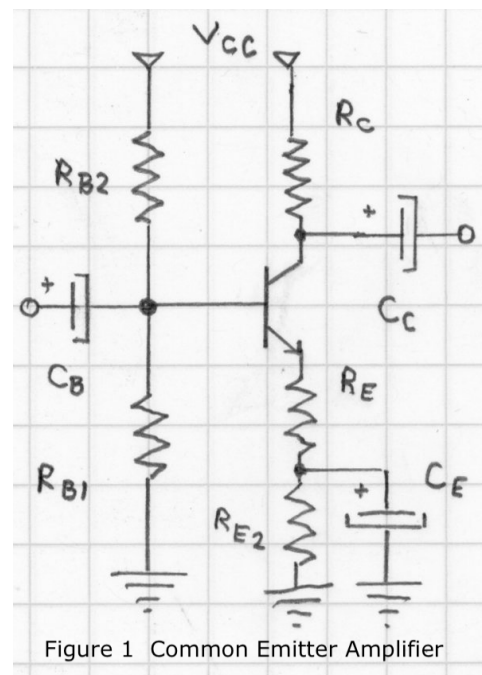


Figure 1 Common Emitter Amplifier

The point is, there's a lot going on. Not mentioned but relevant is, what sort of load should be assumed? If the problem doesn't give a load, should the designer assume $R_{load} = R_c$ (for maximum energy transfer)? Or $R_{load} = \text{infinity}$ (since it wasn't mentioned)? This in turn affects the size of the output signal that can be generated without clipping. What will be done with that issue? Usually you want the Q point selected to give equal swings in both directions.

What about frequency response? (In hindsight, I should have asked specifically for an audio amplifier.) How should the poles at the low end of the pass band be distributed? That may depend upon whether this amplifier will be part of a larger system which uses feedback. What would you assume? What should Chat GPT assume? As for the high end, we probably would not be asking for Common Emitter if we were doing an RF circuit. So, audio would make sense. It wouldn't be control, because the capacitors mentioned block DC.

There is a lot going on in this design problem. How much of that will Chat GPT grasp? How well will it be able to express the design in words? (Beyond the scope I thought of: explaining the design process, but that's what it actually did, at least to some extent.)

Chat GPT's response to the design problem:

To design a capacitor-coupled common-emitter (CE) bipolar amplifier for a 12 V power supply with a 10 kilo-ohm input impedance and maximum gain, you would need to follow these steps:

1. Determine the operating point: Let's assume we choose a quiescent collector current of 1 mA and a collector voltage of 6V. To achieve this, we can use a voltage divider consisting of two resistors to bias the base of the transistor at about 0.7 V, and we can set the emitter resistor to 6.8 kohm.
2. Calculate the required gain: Let's assume we want a voltage gain of 100.
3. Select the coupling capacitor: To ensure that the capacitor does not significantly affect the amplifier's gain at the desired frequency range, we can use a capacitor with a reactance of less than 100 ohm at the lowest frequency of interest. Let's choose a capacitor of 1 uF.
4. Choose the biasing resistors: Using the values from step 1, we can calculate the two biasing resistors. $R_1 = (12 - 0.7) / 10e-6 = 1.193 \text{ MOhm}$ and $R_2 = (12 - 6) / (1 \text{ mA}) = 6 \text{ kohm}$.
5. Calculate the collector load resistor: The collector load resistor should be chosen to maximize the gain of the amplifier. We can calculate this value using the formula $R_c = (V_{cc} - V_{ce}) / I_c$. Plugging in the values, we get $R_c = (12 - 6) / (1 \text{ mA}) = 6 \text{ kohm}$.
6. Build the circuit: Using the values from the previous steps, we can construct the circuit as follows:
 - Connect the 1 uF coupling capacitor between the emitter and the base of the transistor.
 - Connect the 1.193 MOhm and 6 kohm biasing resistors in series between the base of the transistor and the 12 V power supply.
 - Connect the 6.8 kohm emitter resistor between the emitter of the transistor and ground.
 - Connect the 6 kohm collector load resistor between the collector of the transistor and the

12 V power supply.

- Connect the input signal source to the base of the transistor and the output load to the collector of the transistor

Analysis of Chat GPT's design process:

We will do this step by step, assuming that the sequence of text given aligns with the sequence of design steps taken.

1. Determine the operating point: Let's assume we choose a quiescent collector current of 1 mA and a collector voltage of 6V. To achieve this, we can use a voltage divider consisting of two resistors to bias the base of the transistor at about 0.7 V, and we can set the emitter resistor to 6.8 kohm.

The choice that I_C is 1 mA is given no justification at all. It is a "choice", not an assumption. (We read, "we choose a quiescent collector current of 1 mA ...") Why did Chat GPT choose 1mA? That is going to turn out to be wrong. Unless, later, a substantial RE (unbypassed emitter resistor) is used. But that will sacrifice gain. Is that where Chat GPT is going? No. It will be seen later that this choice is a mistake. It is, I will add, a common mistake. Students looking at a design problem on a test or homework will often start out by choosing I_c without having considered the design goals of the problem. Chat GPT seems to be doing exactly that. It doesn't understand how to take constraints into consideration. Or, doesn't seem to.

Right here at the beginning, Chat GPT picks an unorthodox Q point of $V_c=6V$ instead of $V_c=8V$. There is a legitimate reason this might have been done. It sacrifices a bit of bias stability and input impedance to give a larger $V_{cc}-V_C$, allowing RC to be larger and hence yield more gain. Yes, we do want gain. Is that what Chat GPT was thinking?

The assumption that V_{BE} is .7 Volts is the usual assumption for Silicon transistors, which is a reasonable default assumption. Nobody said anything about Germanium or Gallium Arsenide. So that's entirely reasonable.

Then we get, "we can set the emitter resistor to 6.8 kohm." No, we can't. If $I_E = I_C$ (and it's actually going to be larger), and $I_C = 1$ mA, then $V_E \geq 6.8V$, and $V_E > V_C$. That means that the transistor is outside the linear region. Indeed, instead of 6.8V, you will see the transistor's I_C fall below 1 mA until V_C is about the same as V_E and the transistor is "saturated." Apparently Chat GPT doesn't understand what the linear region of a transistor is. Which implies Chat GPT doesn't really know what a transistor is or how it operates. For V_C to be lower than V_E with $I_C=1$ mA implies that the transistor is creating energy. Since there is no energy source, it would be violating the First Law of Thermodynamics, in the sense that you can't just create energy from nothing. Does Chat GPT not understand the First Law of Thermodynamics?

I will give Chat GPT credit for at least choosing a "standard value" for the resistor. 6.8K Ohms is a resistor that you can readily purchase in a variety of tolerances. For this inexpensive circuit, we'd normally pick a cheap 5% resistor, unless it was going to the moon or something exotic like that.

2. Calculate the required gain: Let's assume we want a voltage gain of 100.

Here we have a problem. The problem statement said to maximize the gain, not to choose the gain. So, why? It is wrong to choose a gain of 100 when we might have gotten 150. So, what are we doing? Are we going to do gain = 100 then see if Z_{in} is met, and if so, see if we can then get a higher gain? Or if we fail to get 10K, lower the gain and try to see how Z_{in} comes out? So, yes, this can make sense as part of an iterative approach. We'll withhold judgment for the time being. (An assumption of transistor current gain of 100 would be reasonable. And some sort of assumption about DC and/or AC Beta is necessary. That is missing.)

3. Select the coupling capacitor: To ensure that the capacitor does not significantly affect the amplifier's gain at the desired frequency range, we can use a capacitor with a reactance of less than 100 ohm at the lowest frequency of interest. Let's choose a capacitor of 1 μF .

Why are we choosing a capacitor value when we have not yet finished the bias (DC) part of the design? We don't yet know the "R" part of the RC time constant associated with this capacitor. Furthermore, if an assumption is going to be made, shouldn't it be frequency rather than a capacitor value? Chat GPT says it wants less than 100 Ohms at the frequency of interest, but what does that say what that is when it goes ahead and picks 1 μF ? That's equivalent to saying pick $f_b = 1600 \text{ Hz}$. Since the passband is going to start here or higher, perhaps considerably higher, that cuts off a major part of the audio spectrum. What is Chat GPT thinking!?! Or, is it?

4. Choose the biasing resistors: Using the values from step 1, we can calculate the two biasing resistors. $R_1 = (12 - 0.7) / 10\text{e-}6 = 1.193 \text{ MOhm}$ and $R_2 = (12 - 6) / (1 \text{ mA}) = 6 \text{ kohm}$.

Whoa! It has not yet told us where R_1 and R_2 are connected. Looking at the Voltages implied in the equations to calculate the resistances, R_1 looks like it connects between V_{cc} and a diode to ground (the base-emitter junction?). R_2 appears to go between V_{cc} and ... what? The collector? But, we already had an R_E (or R_{E2}) that was 6.8 K Ohms. Where did 10e-6 come from? Or the 1mA? OK, maybe R_2 is simply the collector resistor. And the lower bias resistor is simply left out. But then R_1 should be from V_{cc} to the base at a Voltage .7V above V_E . At this point we don't know what is going on. Chat GPT hasn't adequately explained what it is telling us. Well, the problem statement didn't require an adequate explanation. The fact we are given this is gratuitous. But, oh my, it looks like this design is heading for trouble. (Chat GPT doesn't seem to know that units should be attached to numbers when doing these kinds of problems.)

I have to take back what I said about crediting Chat GPT for knowing about and using standard values. How do we get a resistor of "1.193 MOhm"? We'd have to use a trim pot and do an alignment. Or make a special order at large expense. Use of a 1.2M Ω resistor would make sense. "6 kohm" isn't a standard 5% (inexpensive) standard value resistor either.

5. Calculate the collector load resistor: The collector load resistor should be chosen to maximize the gain of the amplifier. We can calculate this value using the formula $R_c = (V_{cc} - V_{ce}) / I_c$. Plugging in the values, we get $R_c = (12 - 6) / (1 \text{ mA}) = 6 \text{ kohm}$.

The formula used only works if $V_E = 0\text{V}$, that is, the emitter is at DC ground, not simply an AC ground. This implies that That GPT is trying to design a "grounded emitter" circuit in the

most literal sense rather than “common emitter” in the AC sense. Chat GPT doesn’t really understand what we mean by a “Common Emitter amplifier.” Yet, earlier, it assigned an emitter resistor to have the value 6.8K! What was that about if $R_E=0$? (The $R_E=0$ IS compatible with the bias resistor calculation of 1.2 M Ω in part 4 above though.)

So, I’m confused. If the emitter is really literally grounded, an R_C of about 6K Ohms makes sense, giving equal and maximum swing both up and down. But if $R_E=6.8K$, we are not even in the linear region. What is going on????

Furthermore, what happened to the requirement to have an input impedance of 10K Ohms? Was that achieved? What about the determination that we will get a (Voltage) gain of 100? Was that achieved? No answers are given to these questions. Well, I suppose I didn’t ask. I didn’t say, “Tell me what gain and Z_{in} are actually achieved? Why should I need to do that? But, Chat GPT should have at least taken into account, one way or another, the design constraint given for Input impedance. It didn’t. This is like a student leaving out inconvenient constraints because he wants to do a problem different than the one assigned, and then pretends he misunderstood the assignment.

Finally, we held off judging some design choices hoping that these would be corrected by iterative design. But, there are no iterations. Chat GPT has made one pass, spit out a design, and stopped there as if it were satisfactory, without any analysis to see if the design meets the requirements given.

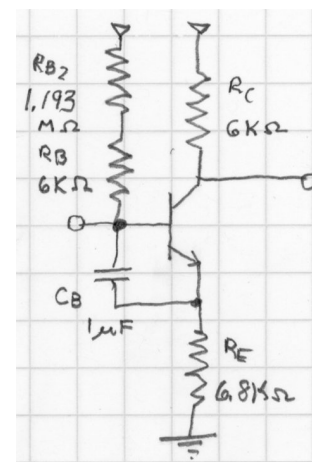
Analysis of the design:

Maybe we can make sense of this if we look at the steps Chat GPT gives us to build this:

- Connect the 1 μ F coupling capacitor between the emitter and the base of the transistor.
- Connect the 1.193 MOhm and 6 kohm biasing resistors in series between the base of the transistor and the 12 V power supply.
- Connect the 6.8 kohm emitter resistor between the emitter of the transistor and ground.
- Connect the 6 kohm collector load resistor between the collector of the transistor and the 12 V power supply.
- Connect the input signal source to the base of the transistor and the output load to the collector of the transistor

Here (at right) is what that looks like. Immediately we see the following faults:

1) There is no input coupling capacitor. The source, if it is it an AC “Voltage source,” will ground the base DC Q point. That will cause, at the least, the transistor to be in cutoff at least half of each AC cycle assuming a sinusoid signal. Maybe that’s OK for a class C amplifier. Did Chat GPT decide arbitrarily to build in Class C? No, because the capacitor between the base and the emitter is supposed to be the coupling capacitor. The fact is, Chat GPT doesn’t understand what a coupling capacitor is or what it is supposed to do!



- 2) The bias resistors are in series. They could be replaced with one resistor.
- 3) The bias resistance is only correct for an emitter Q point Voltage of 0 Volts, but that can't happen with the large RE given.
- 4) The absence of a bias resistor from base to ground, given that the emitter actually was grounded, means that the circuit has an extreme sensitivity to variations in Beta.
- 5) The transistor Voltage gain is crippled by the absence of an emitter bypass capacitor to be less than unity, even if other errors were corrected. That's not "maximum gain."
- 5) If the emitter is assumed to actually be grounded, with the designed Q point $I_c=1\text{mA}$, then the input impedance (for AC) would be about 2.6 K Ohms, well below the specification.
- 6) The capacitor, being where it is, cripples the transistor's AC performance.
- 7) The absence of the output C is perhaps excusable. The coupling capacitor there may be assumed a component of the next stage. But, Chat GPT didn't say so. If it is a resistive load to ground, the usual default for an unknown load, the lack of a capacitor will shift the Collector Q point from what was designed.

So, to summarize, the circuit designed by Chat GPT won't work. Chat GPT makes no effort to assess whether this design works, or whether it meets to performance goals.

Observations:

So, what we see in response to the question is a defective design process that ignores one of the constraints, perhaps the most important one, fails to make needed assumptions, for example transistor Beta, makes a choice for the design gain, which was supposed to be maximized, and seems to not understand the use of capacitors in this circuit. The design not only doesn't work, it doesn't work for numerous reasons. Chat GPT does not check its work to see if what it produced performs at all, much less whether it meets the specified performance goal.

Since one of the concerns about AI is whether it will replace human beings in jobs, the results of this test, at least within the domain of interest here, is a decisive "No." The behavior of Chat GPT seems most comparable to that of an unmotivated student taking an electronics course that he really doesn't care about. He doesn't understand the theory and concepts being taught in the course, so on this test question he just sort of rambles off stuff that doesn't really make sense but might slip past an inattentive grader for at least a few points of partial credit. Is that the kind of student you want to hire for a job doing design work?

The argument can be made that, well, Chat GPT shouldn't be expected to make reasonable assumptions or test its answers or make sure it answers design goals. That the 'bot needs to be given feedback that these various errors were present. I don't know how many such cycles of feedback and iterations on the specification would be necessary. My guess is, quite a lot. The experience would be like a supervisor with an employee who doesn't seem to care whether his work is any good or not. He comes to the supervisor with poor work, as if it is his supervisor's job to give him the motivation, instruction, and precise directions needed for him to

get the job done. If the supervisor has to give all of that help to this one “employee” in order to get satisfactory performance, how many employees can he afford to have working for him, as well as getting his own management job done? Not very many! In most jobs, this guy is going to be out of work shortly.

In a sense, the fundamental problem is this. Chat GPT does not care enough about what it produces to test its answer. In the case of a circuit, that could be done using simulation. OK, maybe Chat GPT doesn't have the tools, the initiative, or the license to do that. But an analytic analysis of the design is still very possible. For example, draw the DC (bias) model of the circuit then use loop or node analysis. These are algorithms. Chat GPT should be able to do that without being told. Once the Q point is known from DC analysis, there are straightforward means for doing AC analysis to find the mid-band Voltage gain, and the input impedance, and the frequency response. None of that was done, even just to find if the impedance specification was met, or whether the chosen gain goal was met.

It could be argued that Chat GPT is only in demo mode, and that it is not required to take this test problem seriously. After all, the problem was not fully specified, requiring assumptions and perhaps even guesses about the application. Why should Chat GPT take this any more seriously than the questioner? If asked a BS question, why should it not give a BS answer? Which is what it gave. However, most software packages that are intended to demo their capabilities do not limit the functionality that is intended to impress potential users with their amazing capabilities. Instead, the ability to keep results, say by disabling save or print features, or by limiting the problem size, are more typical. I find it difficult to believe that Chat GPT has been artificially crippled in this demo mode because the user isn't paying. I'd guess that in “full” or “paid” mode it is going to behave similarly.

Conclusions:

Notwithstanding all that has been written above, Chat GPT is an amazing achievement. It is capable of doing things that were not so routinely done before. The deficiencies noted are in comparison to what a human with specialized education in this domain might be expected to do.

So, the way I see it, Chat GPT is quite an achievement! It's not that it isn't, but the standard by which we are comparing it, human performance, is actually more amazingly competent than we actually thought! Much of what is done in textbooks, even, doesn't really dwell on human things that we usually take for granted. Things that we think of as routinely as ethical behavior (telling the truth), professionalism (caring about our product), or understanding the problem in context. The textbooks instead focus on what is hard for humans: the theory, the nomenclature, the math, and the procedures for various analytic tasks.

Humans have typically been checking their work, at least work they cared about, since elementary school. I recall that back then we would add back a subtracted number to see if that comes out with the number we started with. That shouldn't need to be taught in later courses, and often isn't. It's taken for granted that one should re-read work to check for typos and consistency and errors of various sorts. Chat GPT does not have that deep background in human culture, apparently. Why should it care if its results are right or wrong? What are the consequences of being wrong? Those are not understandable concepts, it would seem.

So, no, I'm not worried. There will be jobs where uncaring and lazy employees can profitably be replaced with Chat GPT or something similar. But if you want work from someone

who cares, and mentally models you and your responses well enough to understand the real requirement to seek the company's greater good, that person isn't going to be replaced by the ilk of Chat GPT.

Consider a telephone operator or receptionist. Right now such people have been replaced by voice automated responses. Those automated systems are notorious for sending people on circular response chains, being unable to understand the need, and causing endless and often repetitious delay. They are not nearly as effective as a human being who cares about clients and seeks the good of the client and the company. But, they are cheaper. Maybe the business can afford to alienate and frustrate people. That seems particularly true of medical providers. In those cases, Chat GPT type 'bots can perhaps do a better job. But they likely won't be able to do what a caring human being is capable of anytime soon. Because they don't have the basic understanding of the human condition the way a person does.

In technical fields, in this case electronics, it would seem that the same is true. Chat GPT just plain does not understand circuits. It's as if it tries to do things entirely by rote memorization of factoids without an understanding of how they fit together. Maybe that can be fixed. But the ability to make needed reasonable assumptions also seems to be missing. Chat GPT, like many students who have not been "on the job" yet, seems to have a hard time understanding the problem context, and why they should care. My guess is, that won't change anytime soon.

Reflection:

This has been an interesting exercise. I'm impressed with Chat GPT. I expect to continue playing with it when I can find the time, and may find that in other domains it does better. We'll see. It is a lofty and ambitious goal to aspire to represent human intelligence, or create something that gets even a little bit closer to that goal. AI researchers have been working on that for a long time. But, it seems that for every advance, the realization dawns that what the human mind is capable of is more wondrous and subtle than we realized. The things we take for granted turn out to be amazingly hard. Like caring, and what that implies. It's a fascinating journey.